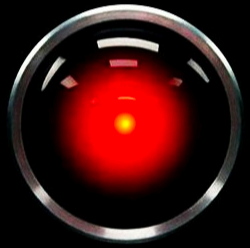
Brief History of Artificial Intelligence:
From Societal and Political to Scientific Impacts…



Miguel Couceiro
(Many thanks to Mário Figueiredo!)

I'm sorry Dave, I'm afraid I can't do that

IST LISBOA inesc id lisboa instituto de telecomunicações ellis European Laboratory for Learning and Intelligent Systems

Artificial Intelligence (AI)

- Simulation of "intelligence" by machines
- Behavior of an artificial system that humans see as intelligent

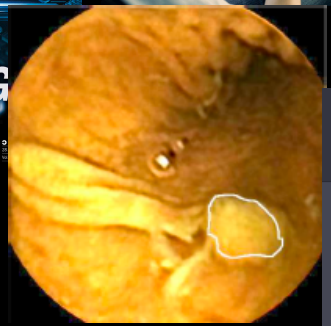












MC Give a concise definition of Artificial intelligence

C Artificial intelligence is the field of computer science concerned with building systems that perform tasks normally requiring human intelligence. Examples include perceiving, reasoning, learning, understanding language and making decisions.



VOL. LIX. No. 236.] [October, 1950

MIND

A QUARTERLY REVIEW

OF

PSYCHOLOGY AND PHILOSOPHY

I.—COMPUTING MACHINERY AND INTELLIGENCE

BY A. M. TURING

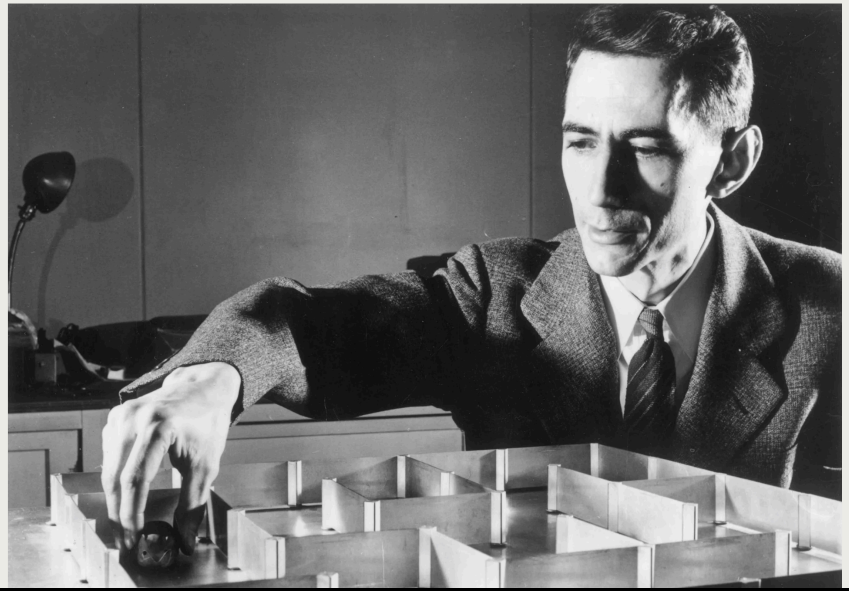
1. *The Imitation Game.*

I PROPOSE to consider the question, 'Can machines think?' This should begin with definitions of the meaning of the terms 'machine' and 'think'. The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words 'machine' and 'think' are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, 'Can machines think?' is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.

The new form of the problem can be described in terms of a game which we call the 'imitation game'. It is played with three people, a man (A), a woman (B), and an interrogator (C) who may be of either sex. The interrogator stays in a room apart from the other two. The object of the game for the interrogator is to determine which of the other two is the man and which is

How Claude Shannon Helped Kick-start Machine Learning › The "father of information theory" also paved the way for AI

BY RODNEY BROOKS | 25 JAN 2022 | 3 MIN READ |



PRESENTATION OF A MAZE-SOLVING MACHINE

Cybernetics - Trans. of the 8th Conference 1952, Josiah Macy, Jr., Foundation, 565 Park Ave., N.Y. 21, N.Y.

CLAUDE SHANNON
Bell Laboratories, Murray Hill, N. J.

THIS IS A maze-solving machine that is capable of solving a maze by trial-and-error means, of remembering the solution, and also of forgetting it in case the situation changes and the solution is no longer applicable. I think this machine may be of interest in view of its connection with the problems of trial-and-error learning, forgetting and feedback systems.

As you can see (Figure 8), there is a maze on the top panel of the machine which has a range of 5 x 5 squares. The maze can be changed in any desired manner by rearranging the partitions between the twenty-five squares. In the maze there is a sensing finger, which can feel the partitions of the maze as it comes against them. This finger is moved by two motors, an east-west motor and a north-south motor. The problem facing the machine is to move the finger through the maze to the goal. The goal is mounted on a pin which can be slipped into a jack in any of the twenty-five squares. Thus you can change the problem any way you choose, within the limits of the 5 x 5 maze. I will turn it on so you can see it, in the first place, trying to solve the maze. When the machine was turned off, the relays essentially forgot everything they knew, so that they are now starting afresh, with no knowledge of the maze.

Savage: Does than mean they are in a neutral position, neither to the right nor the left?

Shannon: They are in a kind of nominal position. It isn't really a neutral position but a meaningless one.

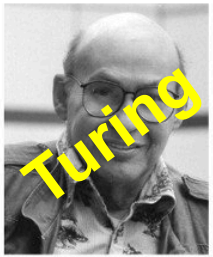
You see the finger now exploring the maze, hunting for the goal. When it reaches the center of a square, the machine makes a new decision as to the next direction to try. If the finger hits a partition, the motors reverse, taking the finger back to the center of the square, where a new direction is chosen. The choices are based on previous knowledge and according to a certain strategy, which is a bit complicated.

Pitts: It is a fixed strategy? It is not a randomization?

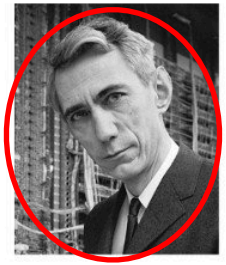
Conferência de Dartmouth: 1956

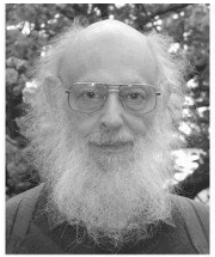


John MacCarthy



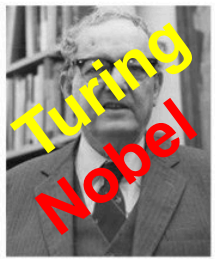
Marvin Minsky


Claude Shannon


Ray Solomonoff



Alan Newell

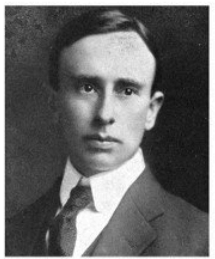


Herbert Simon


Arthur Samuel


Oliver Selfridge


Nathaniel Rochester


Trenchard More

McCarthy coined the term “Artificial Intelligence” over “cybernetics” or “automata theory” and this event gathered 10 people that set AI as a new, independent discipline to address language, abstraction, problem solving and self-improvement… (Too optimistic!!!)

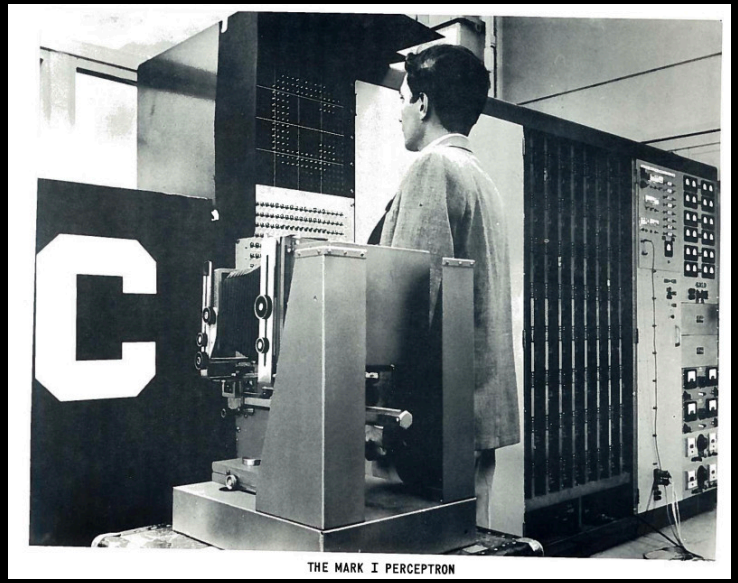
Machine learning

"making computer able to learn without being explicitly programmed"

Arthur Samuel, 1959

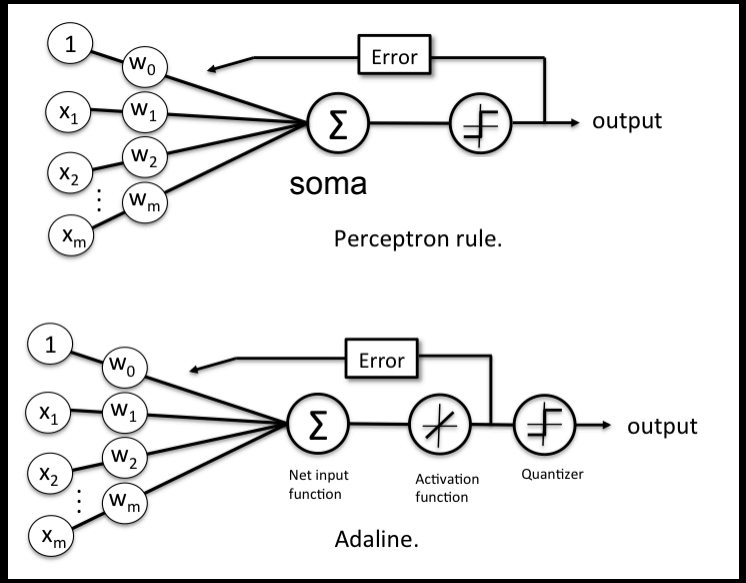


Machine learning



Frank Rosenblatt, 1957
perceptron

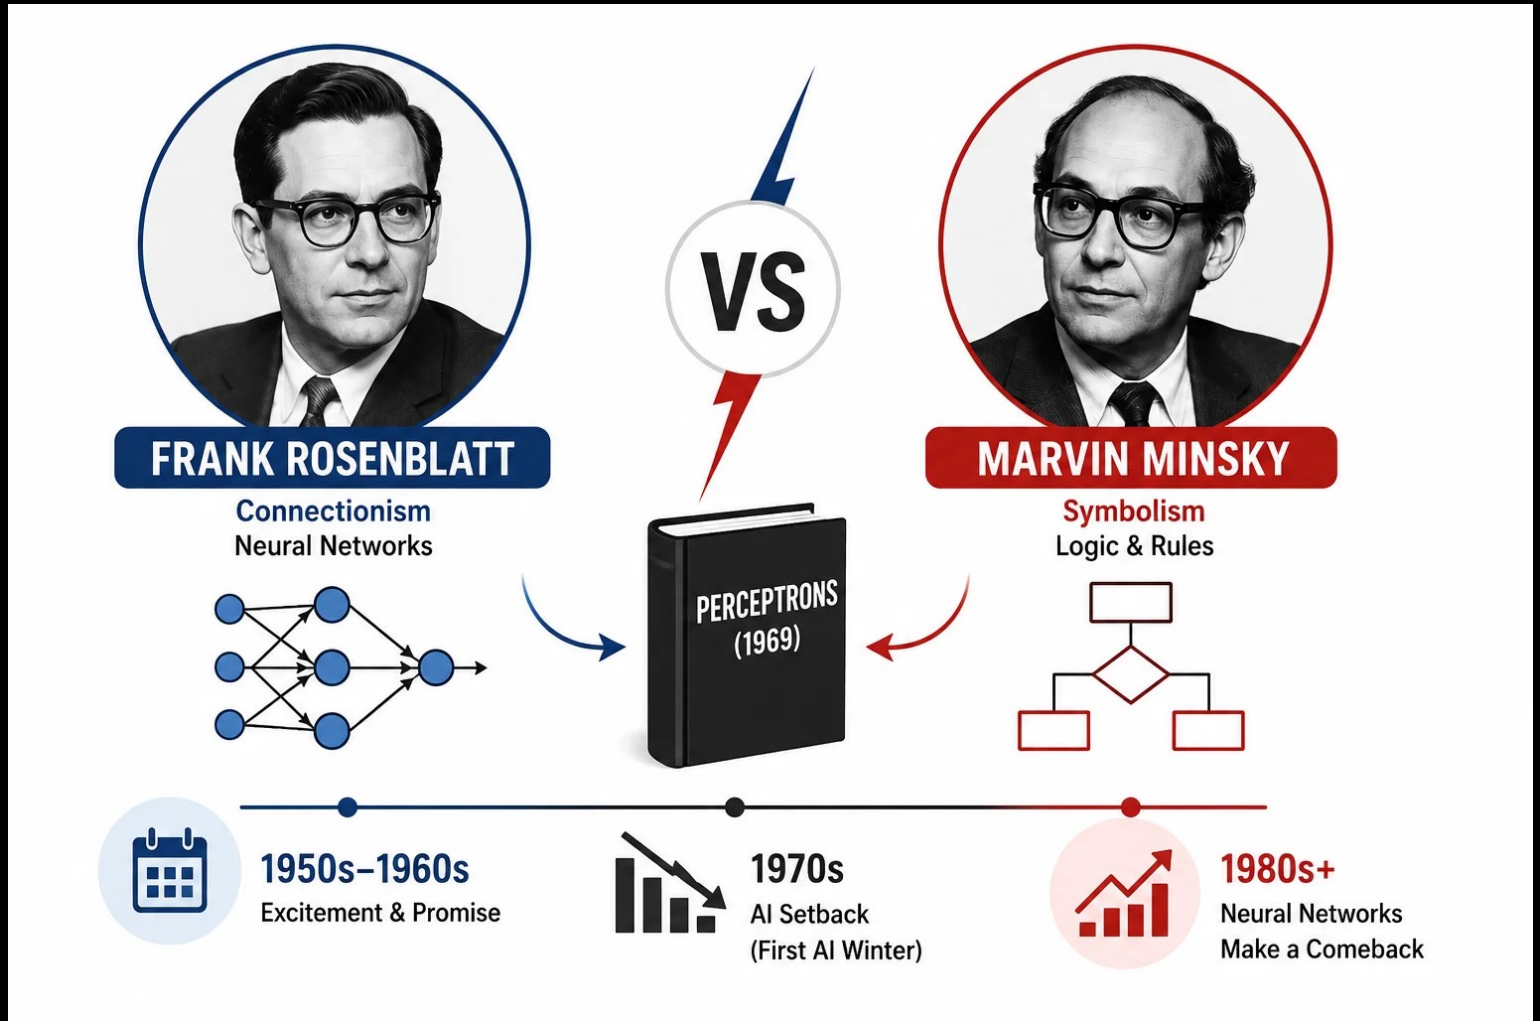
Artificial neurons capable of
"learning" by error feedback.



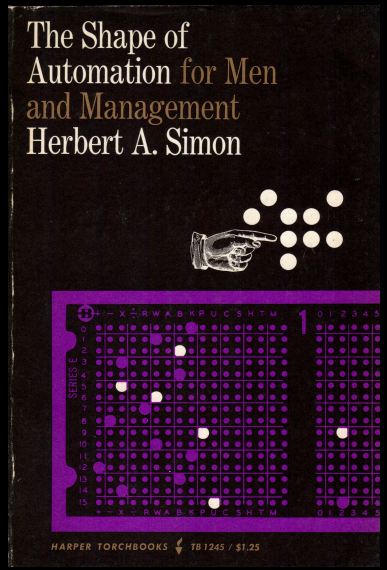
beginnings of artificial neural networks;
beginnings of machine learning
... still the essence of modern systems.

The AI Battle: Marvin Minsky (CMU) vs Frank Rosenblatt (MIT)

(by Abdul Hafeez)



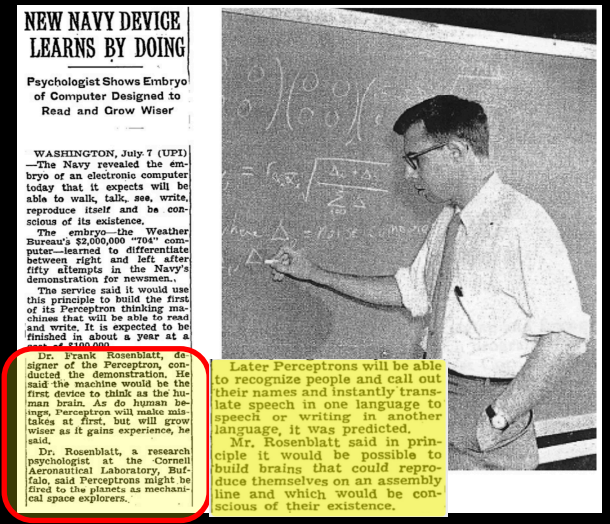
Premature optimism (on all fronts)



"in 20 years, machines will do any human task" 1965



"within a generation (...) the problem of creating AI will be solved" 1967


NEW NAVY DEVICE LEARNS BY DOING

Psychologist Shows Embryo of Computer Designed to Read and Grow Wiser

WASHINGTON, July 7 (UPI) —The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

The embryo—the Weather Bureau's $2,000,000 "704" computer—learned to differentiate between right and left after fifty attempts in the Navy's demonstration for newsmen.

The service said it would use this principle to build the first of its Perceptron thinking machines that will be able to read and write. It is expected to be finished in about a year at a cost of $100,000.

Dr. Frank Rosenblatt, designer of the Perceptron, conducted the demonstration. He said the machine would be the first device to think as the human brain. As do human beings, Perceptron will make mistakes at first, but will grow wiser as it gains experience, he said.

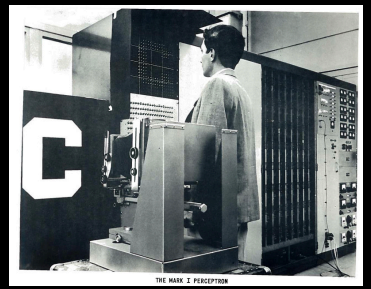
Dr. Rosenblatt, a research psychologist at the Cornell Aeronautical Laboratory, Buffalo, said Perceptrons might be fired to the planets as mechanical space explorers.

Later Perceptrons will be able to recognize people and call out their names and instantly translate speech in one language to speech or writing in another language, it was predicted.

Mr. Rosenblatt said in principle it would be possible to build brains that could reproduce themselves on an assembly line and which would be conscious of their existence.

"in principle it would be possible to build brains that could reproduce themselves ... and which would be conscious of their existence." 1967

"Neural" networks: 3 first decades



Frank Rosenblatt, perceptron, 1957

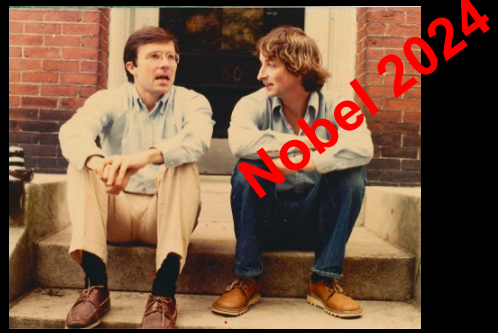


Hopfield networks, 1982
Prior work by Amari, 1972



Rumelhart, Hinton, Williams, backpropagation, 1986

Prior work by Linnainmaa (1970, 1976),
Werbos (1974), LeCun (1985)

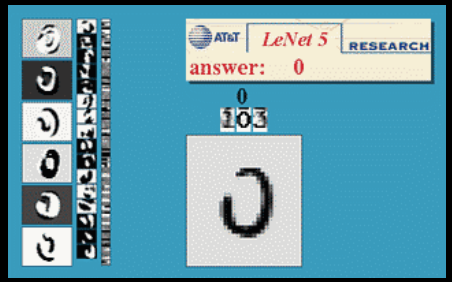


Sejnowski & Hinton, Boltzman machines, 1983.
Prior work by Ivakhnenko & Lapa, 1965, Amari (1967)



LeCun, deep convolutional networks, 1989 (inspired by Hubel & Wiesel),
Prior work by Fukushima (1979), …

1998



Supervised learning: classification

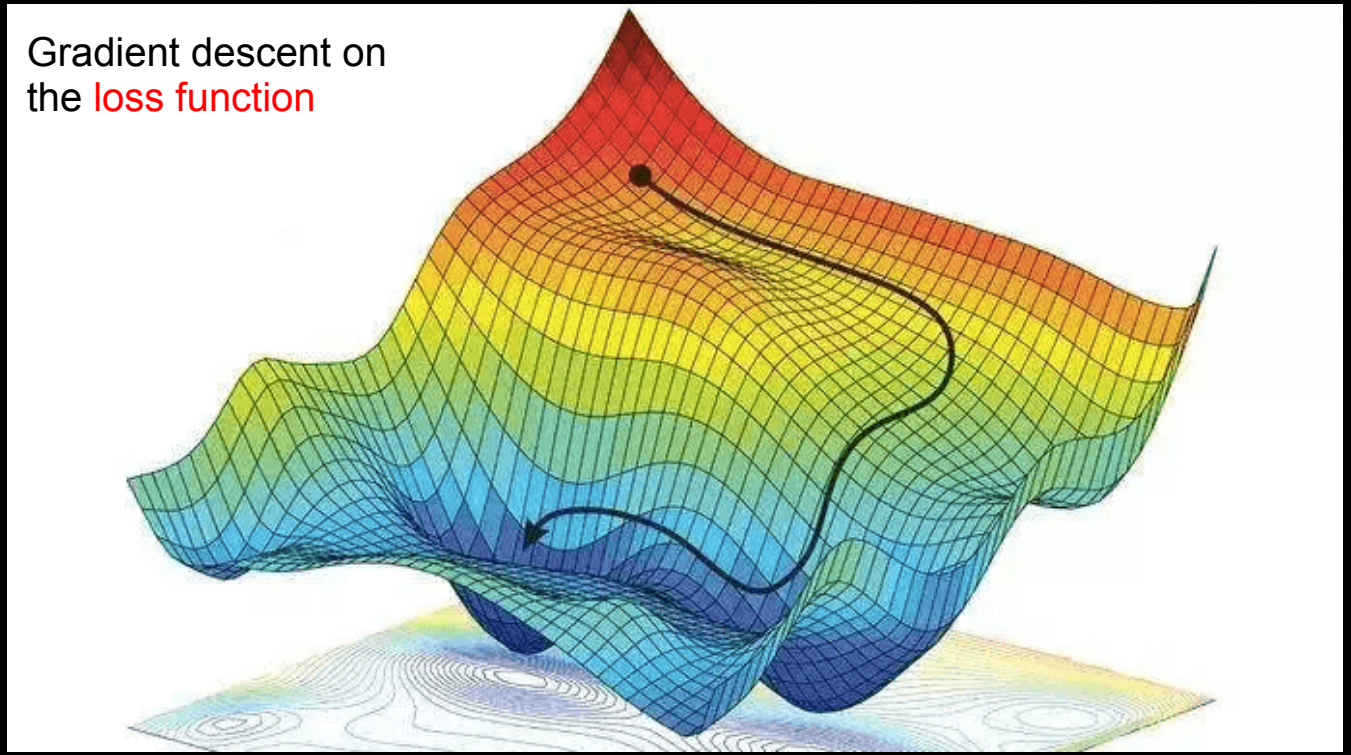




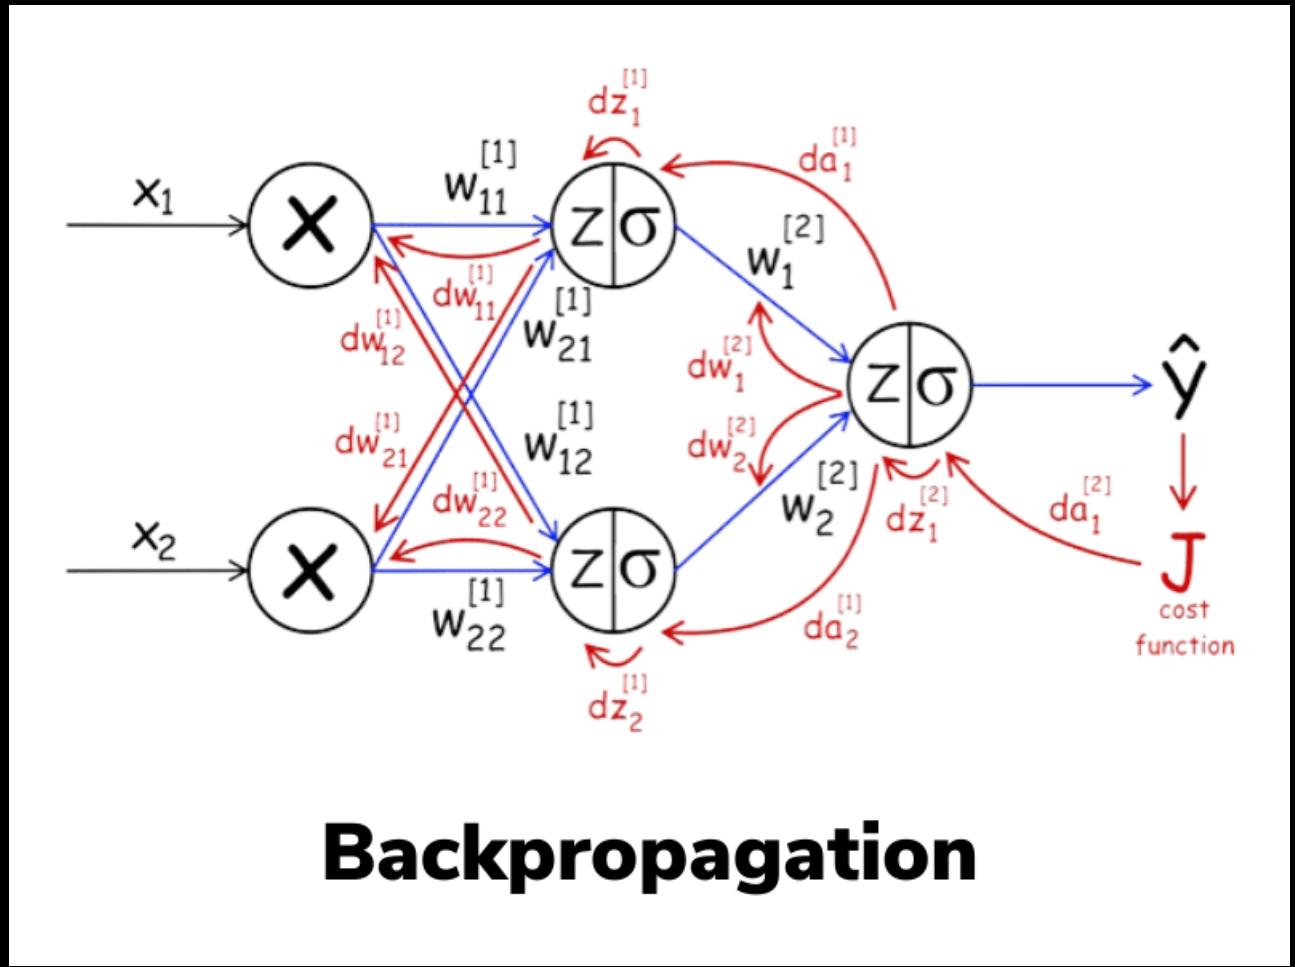
dog

What is a learning algorithm?

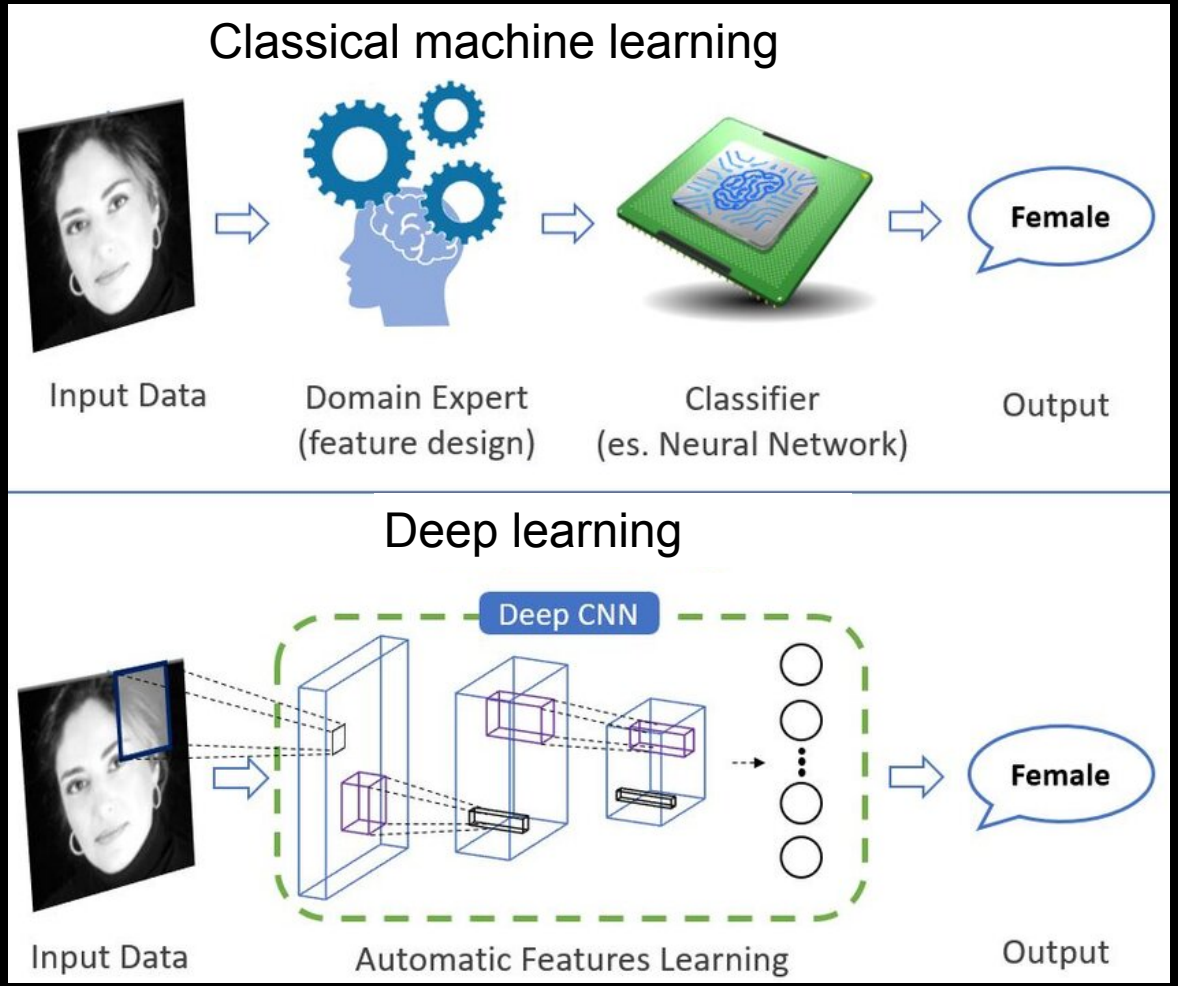
Gradient descent on
the loss function



The secret sauces of deep learning: gradient descent + backpropagation

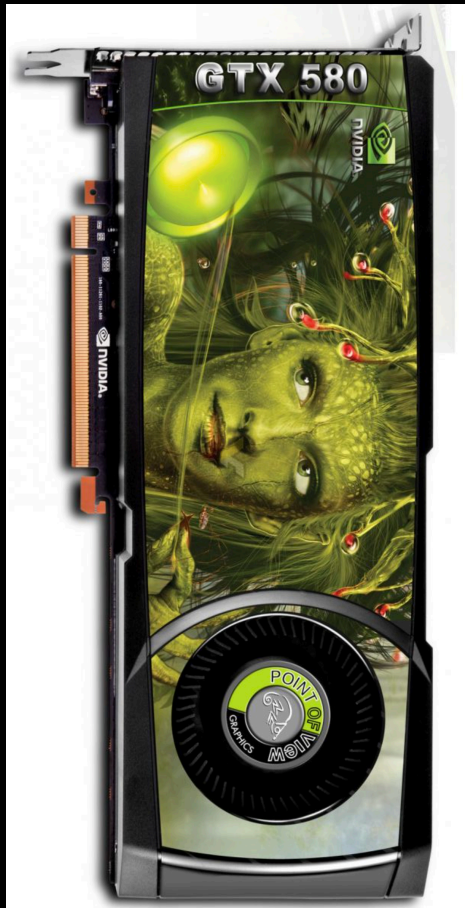


“end-to-end learning”: can be computationally very expensive…



Graphics processing unit (GPU)

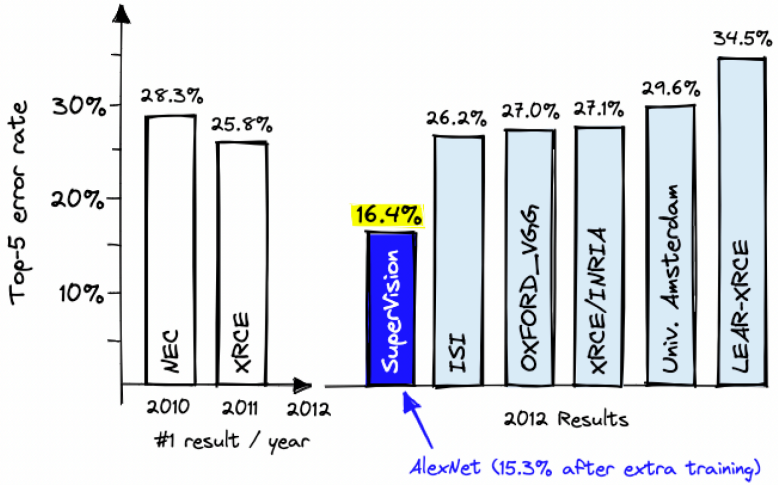




ImageNet 2012

ImageNet (2009): 1.2 million images,
1000 categories

ImageNet Large Scale Visual Recognition Challenge



ImageNet Classification with Deep Convolutional Neural Networks

Alex Krizhevsky
University of Toronto
kriz@cs.utoronto.ca

Ilya Sutskever
University of Toronto
ilya@cs.utoronto.ca

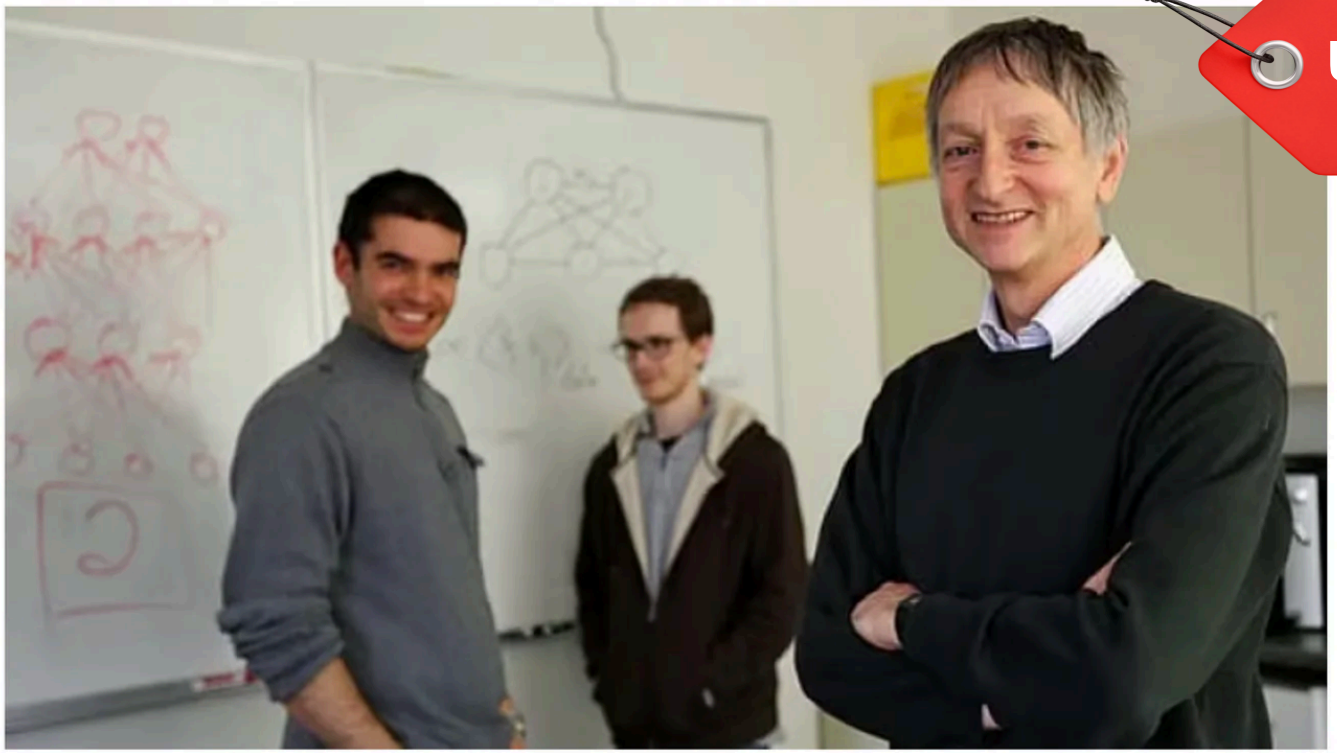
Geoffrey E. Hinton
University of Toronto
hinton@cs.utoronto.ca

Abstract

We trained a large, deep convolutional neural network to classify the 1.2 million high-resolution images in the ImageNet LSVRC-2010 contest into the 1000 different classes. On the test data, we achieved top-1 and top-5 error rates of 37.5% and 17.0% which is considerably better than the previous state-of-the-art. The neural network, which has 60 million parameters and 650,000 neurons, consists of five convolutional layers, some of which are followed by max-pooling layers, and three fully-connected layers with a final 1000-way softmax. To make training faster, we used non-saturating neurons and a very efficient GPU implementation of the convolution operation. To reduce overfitting in the fully-connected layers we employed a recently-developed regularization method called "dropout" that proved to be very effective. We also entered a variant of this model in the ILSVRC-2012 competition and achieved a winning top-5 test error rate of 15.3%, compared to 26.2% achieved by the second-best entry.

Google buys University of Toronto startup

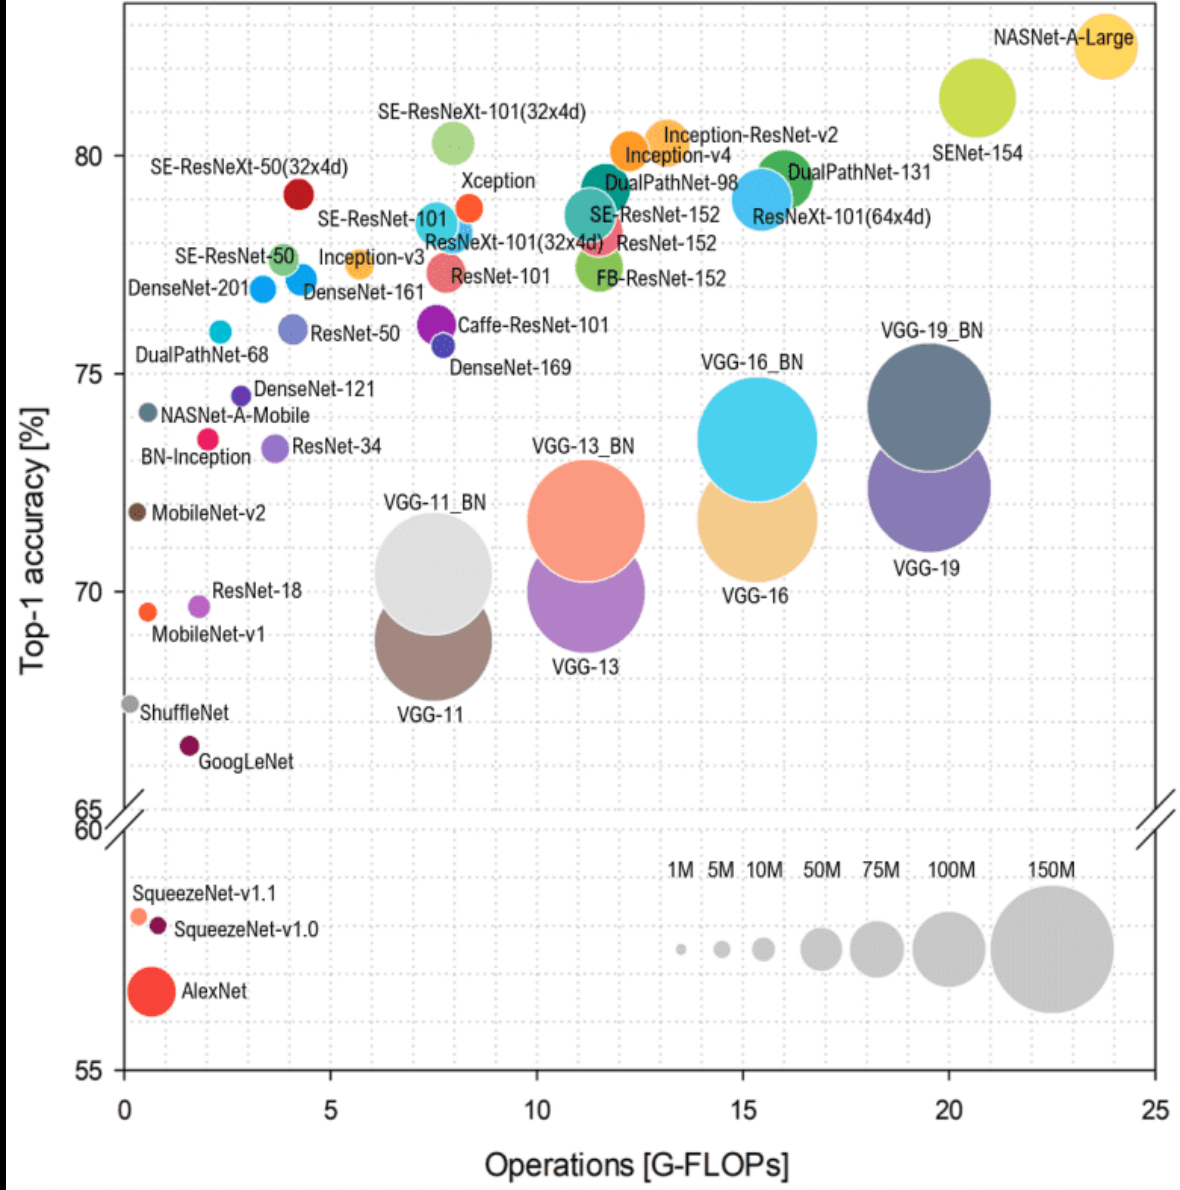
CBC News · Posted: Mar 13, 2013 4:41 PM EDT | Last Updated: March 13, 2013

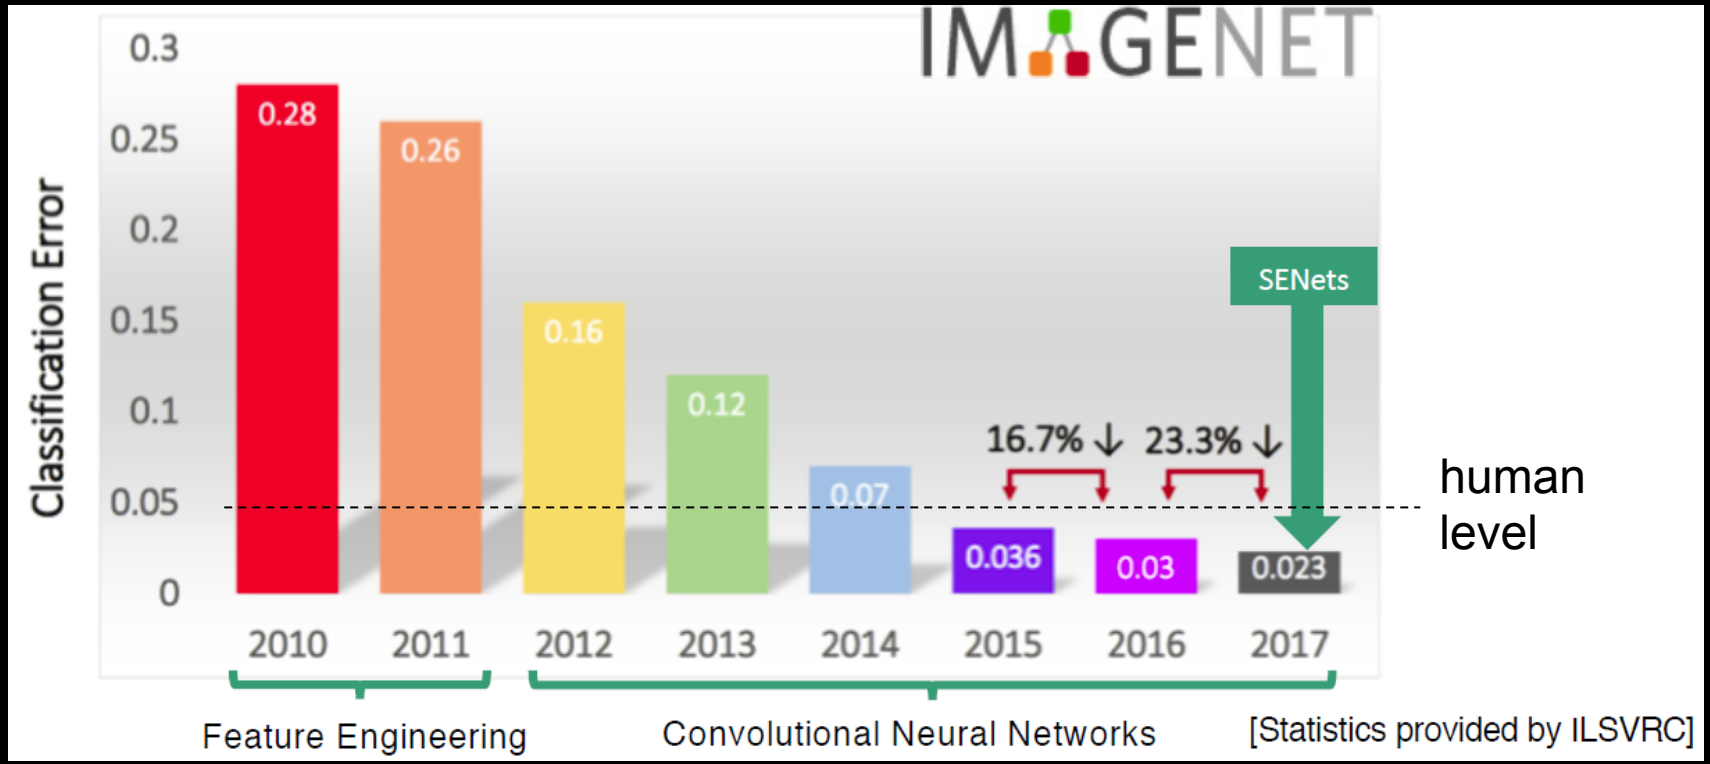


Google has acquired DNNresearch, a University of Toronto startup that studies neural networks. The one-year-old company is launched by computer science professor Geoffrey Hinton (right) and two of his graduate students, Alex Krizhevsky and Ilya Sutskever (left). (The University of Toronto)

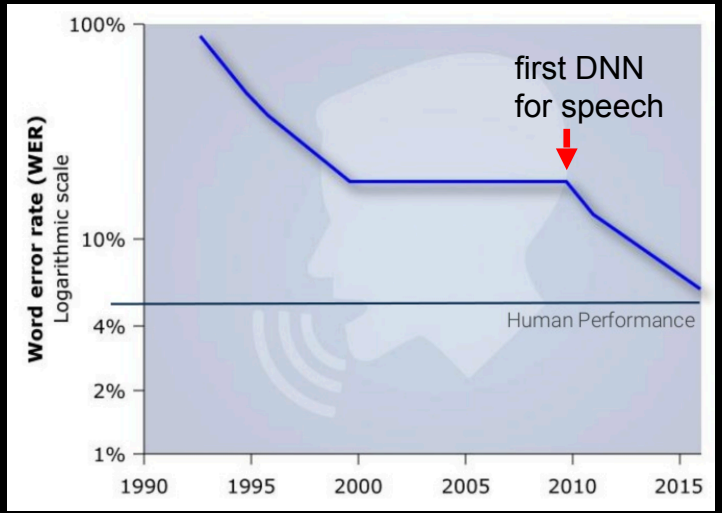
2012-2019

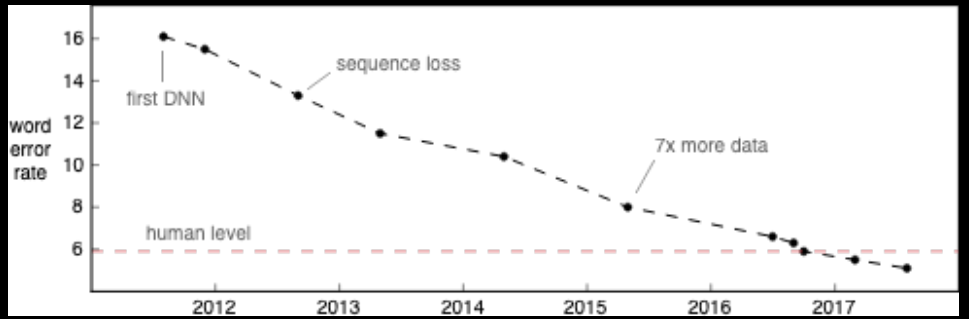
100s of millions of parameters




IMAGENET
Classification Error
0.3
0.25
0.2
0.15
0.1
0.05
0
0.28
0.26
0.16
0.12
0.07
0.036
0.03
0.023
SENets
16.7% ↓
23.3% ↓
human level
2010
2011
2012
2013
2014
2015
2016
2017
Feature Engineering
Convolutional Neural Networks
[Statistics provided by ILSVRC]

Speech recognition





There are no artificial neural networks!

```python
import torch
import torch.nn as nn

# Define the neural network
class Net(nn.Module):
    def __init__(self):
        super(Net, self).__init__()
        self.conv1 = nn.Conv2d(3, 6, 5)
        self.pool = nn.MaxPool2d(2, 2)
        self.conv2 = nn.Conv2d(6, 16, 5)
        self.fc1 = nn.Linear(16 * 5 * 5, 120)
        self.fc2 = nn.Linear(120, 84)
        self.fc3 = nn.Linear(84, 10)

    def forward(self, x):
        x = self.pool(F.relu(self.conv1(x)))
        x = self.pool(F.relu(self.conv2(x)))
        x = x.view(-1, 16 * 5 * 5)
        x = F.relu(self.fc1(x))
        x = F.relu(self.fc2(x))
        x = self.fc3(x)
        return x

net = Net()
```

```python
# Tensorflow Basics on NeuralNetwork with hidden layer
# https://github.com/skcript/tensorflow-resources

import tensorflow as tf
import numpy as np

INPUT_NEURONS = 2
HIDDEN_LAYER_1_NEURONS = 4
OUTPUT_NEURONS = 1

x_data = np.array([[0, 0], [0, 1], [1, 1]])
y_data = np.array([0, 1, 1])

x = tf.placeholder(tf.float32)
y = tf.placeholder(tf.float32)

h1_layer = {
  'weight': tf.Variable(np.random.rand(INPUT_NEURONS, HIDDEN_LAYER_1_NEURONS).astype(np.float32)),
  'bias': tf.constant(0.0)
}
output_layer = {
  'weight': tf.Variable(np.random.rand(HIDDEN_LAYER_1_NEURONS, OUTPUT_NEURONS).astype(np.float32)),
  'bias': tf.constant(0.0)
}


h1 = tf.matmul(x, h1_layer['weight']) + h1_layer['bias'])
h1 = tf.sigmoid(h1)

predict = tf.matmul(h1, output_layer['weight']) + output_layer['bias'])

loss = tf.reduce_mean(tf.square(predict - y))
optimizer = tf.train.GradientDescentOptimizer(0.00001).minimize(loss)

sess = tf.Session()
init = tf.initialize_all_variables()
sess.run(init)

for i in range(50000):
        if i % 5000 == 0:
            print 'Loss: ', sess.run(loss, {x: x_data, y: y_data})
        sess.run(optimizer, {x: x_data, y: y_data})

print sess.run(predict, {x: [[1, 1]] })
```

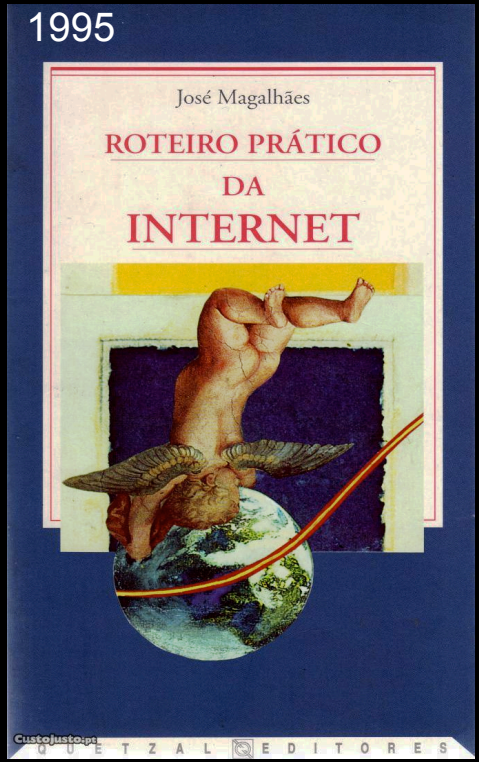
Only 30 years ago...



1994

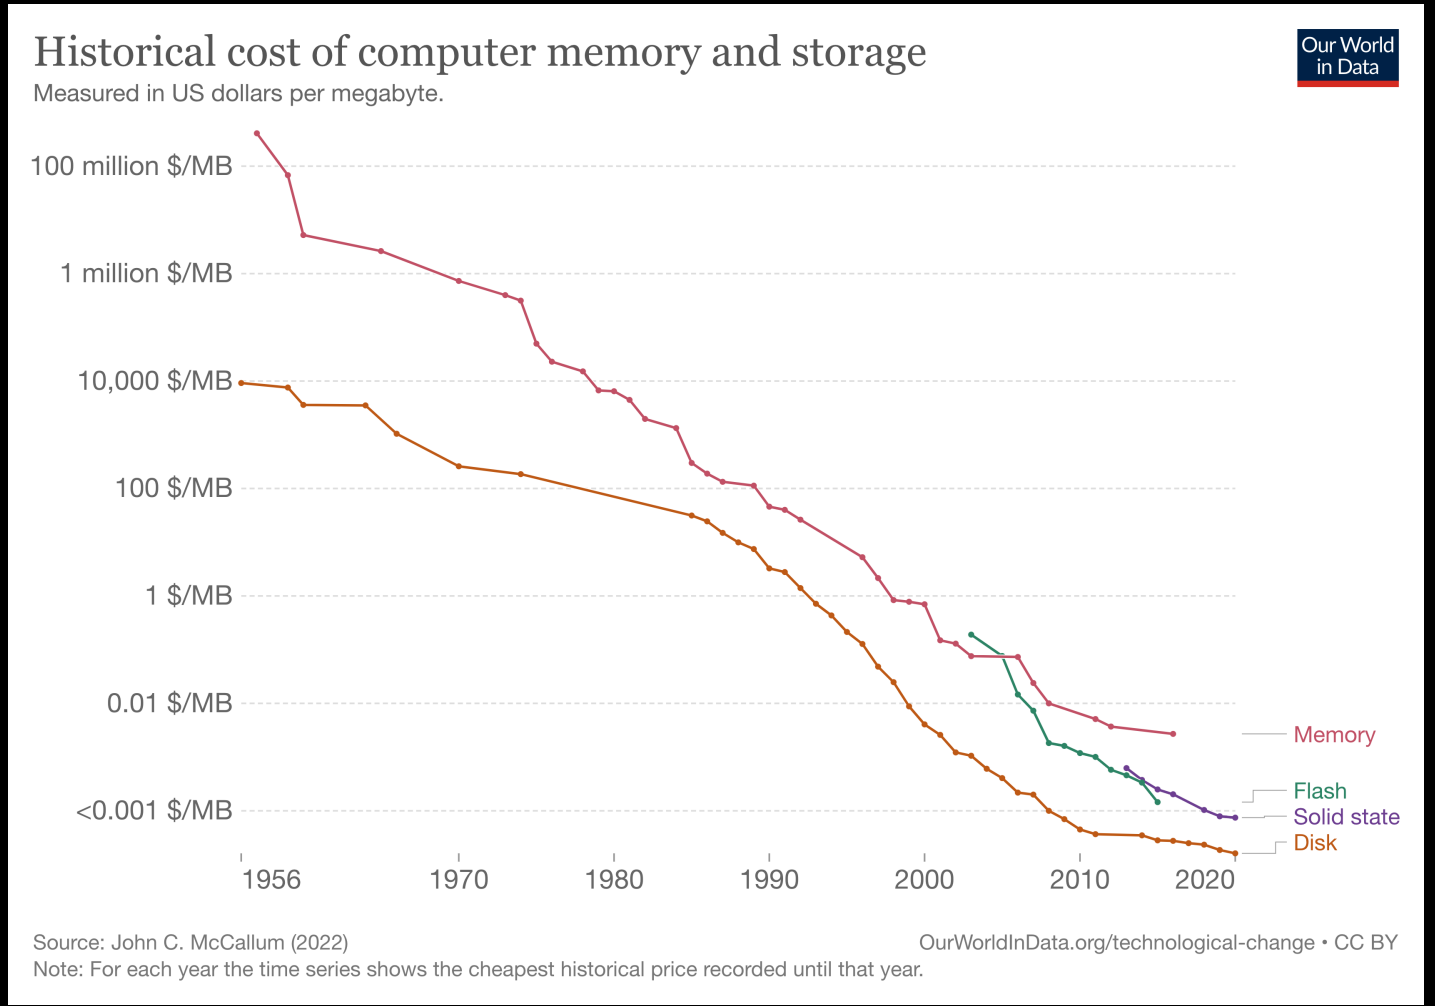


1995

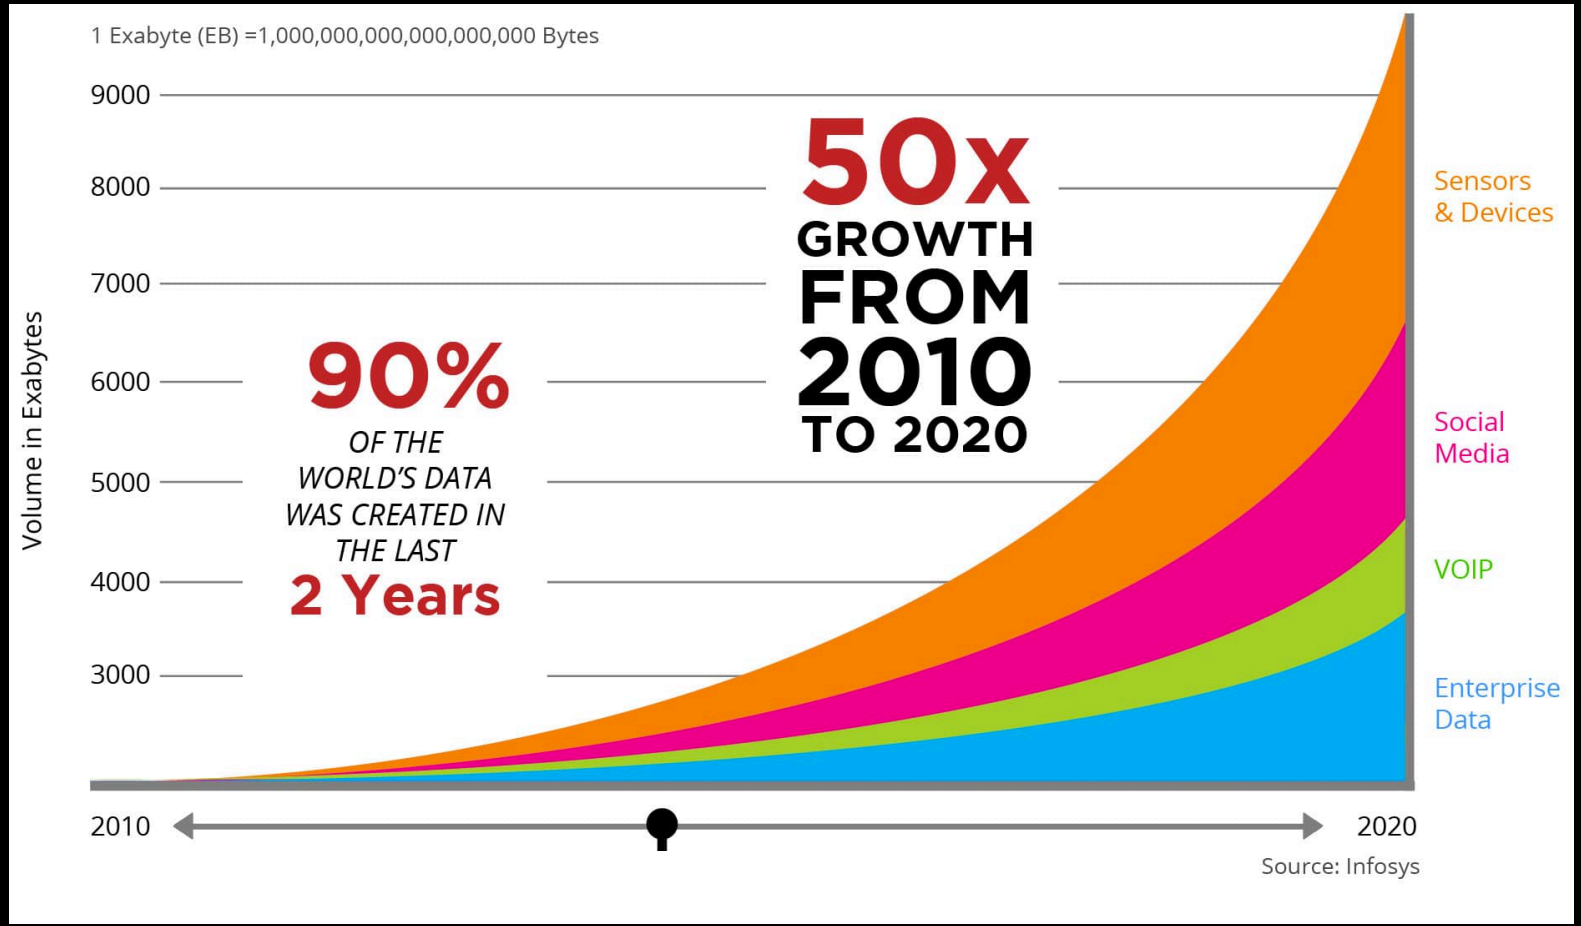


Internet was small:
you could have a “guide”

Technological revolutions



Technological revolutions

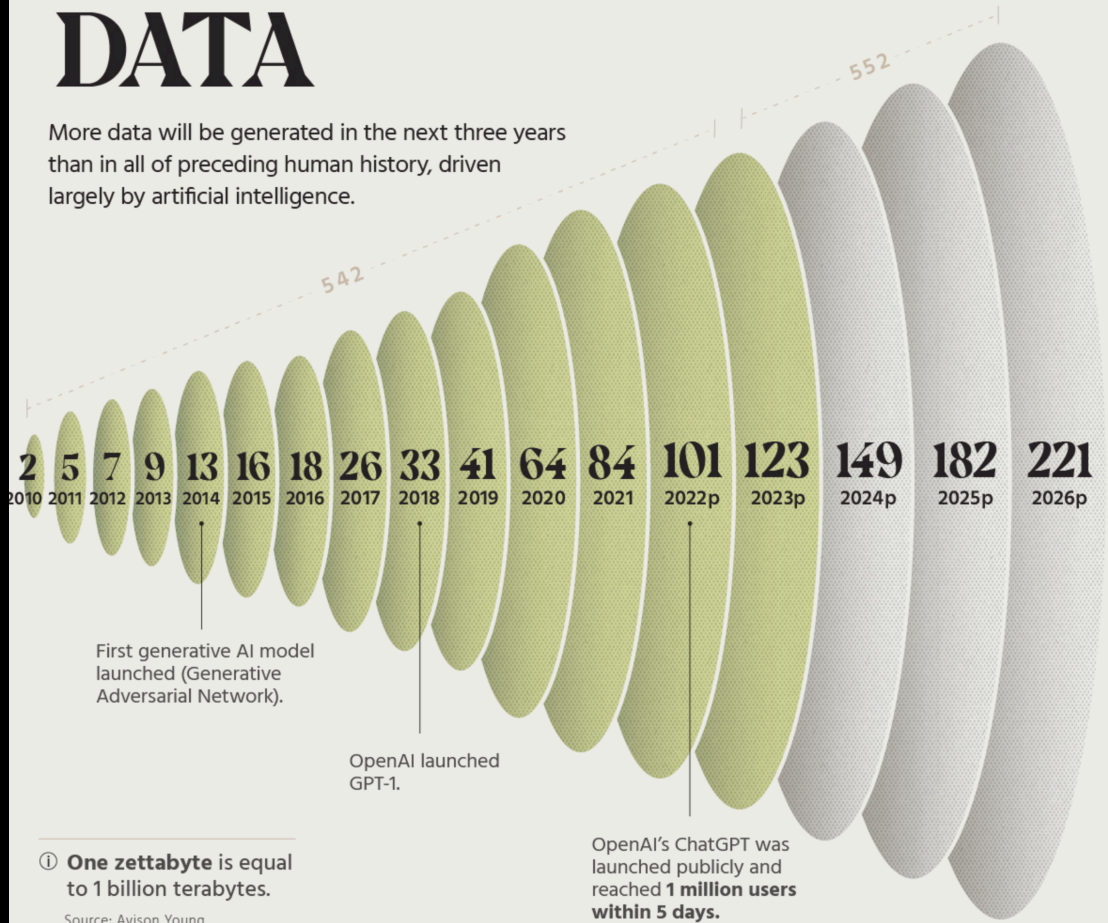


VISUALIZED:

ALL THE WORLD'S DATA

hinrich foundation
advancing sustainable global trade

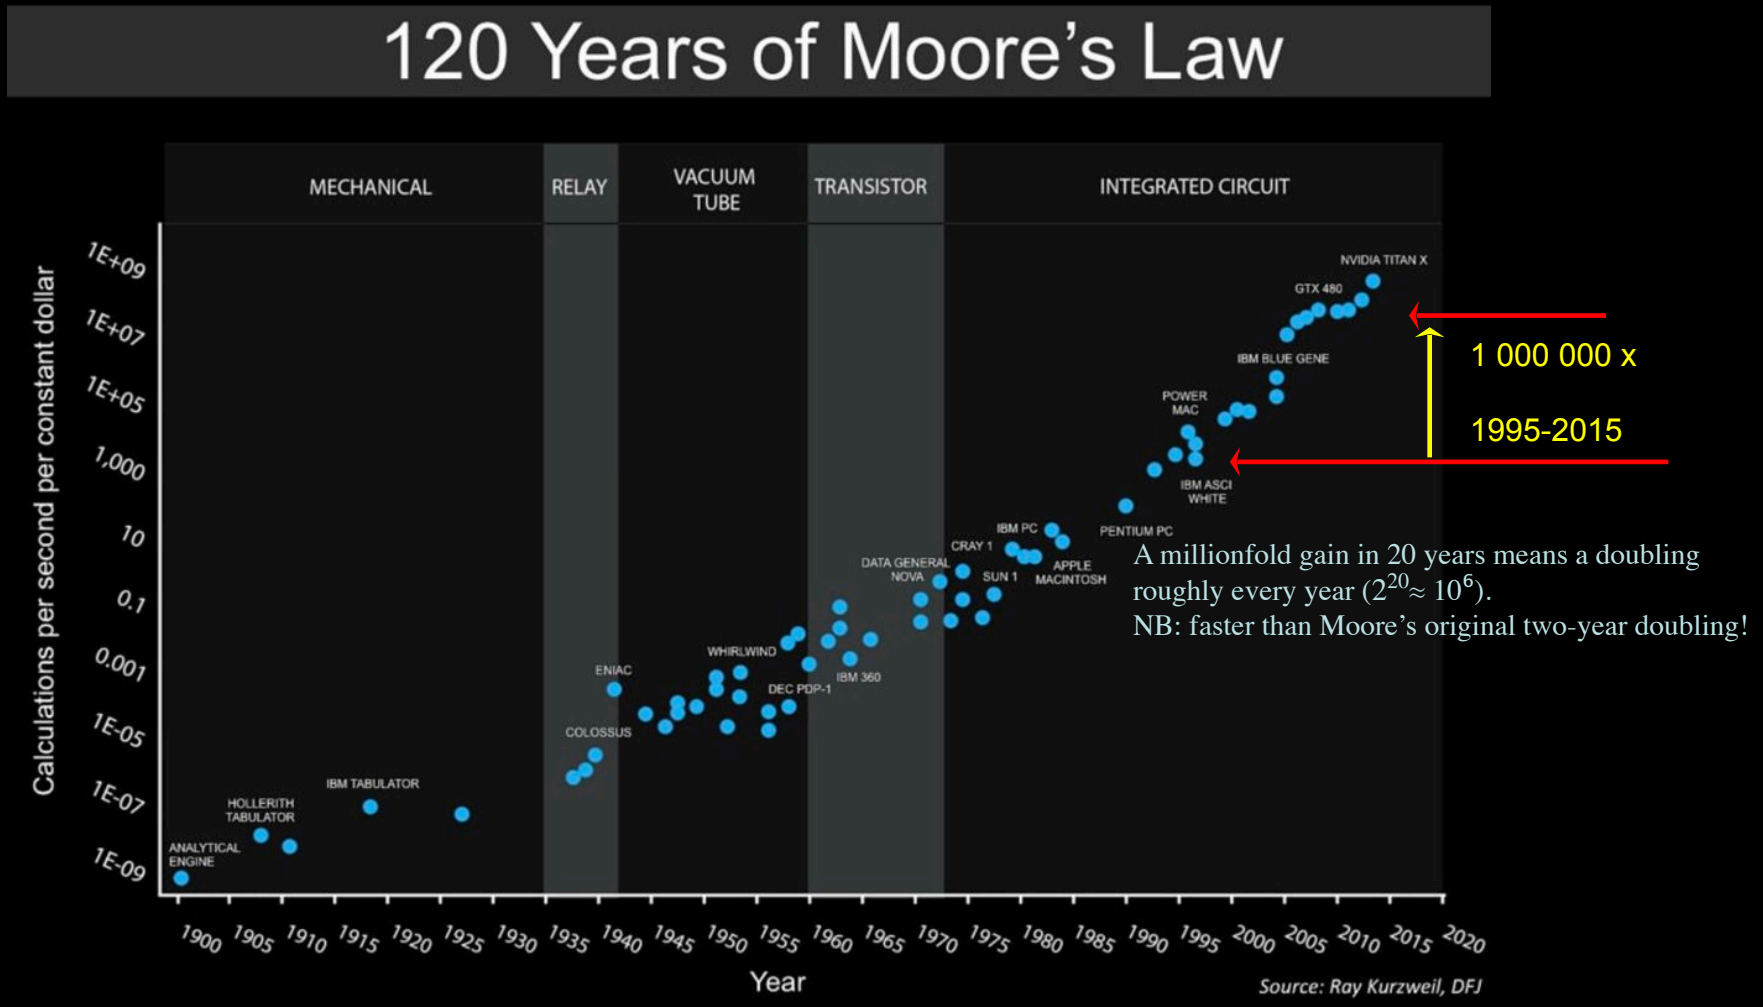
More data will be generated in the next three years than in all of preceding human history, driven largely by artificial intelligence.



ⓘ **One zettabyte** is equal to 1 billion terabytes.

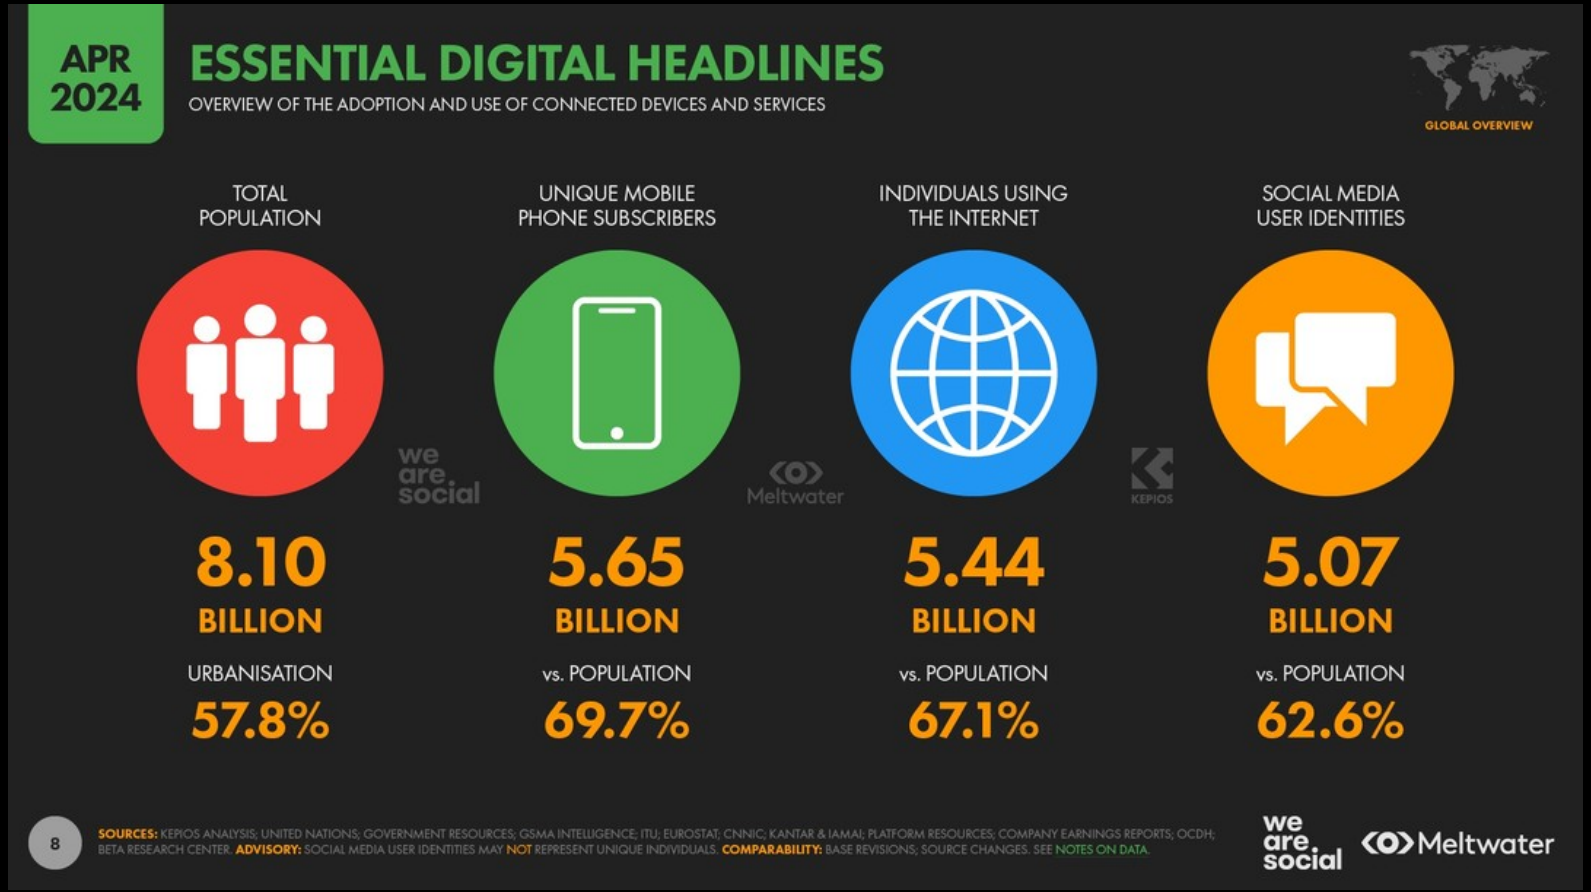
Source: Avison Young

Technological revolutions



A millionfold gain in 20 years means a doubling roughly every year ($2^{20} \approx 10^6$).
NB: faster than Moore's original two-year doubling!

Technological revolutions



Access and generation/production of data:
video, photos, location, markets, social media, health, science, …

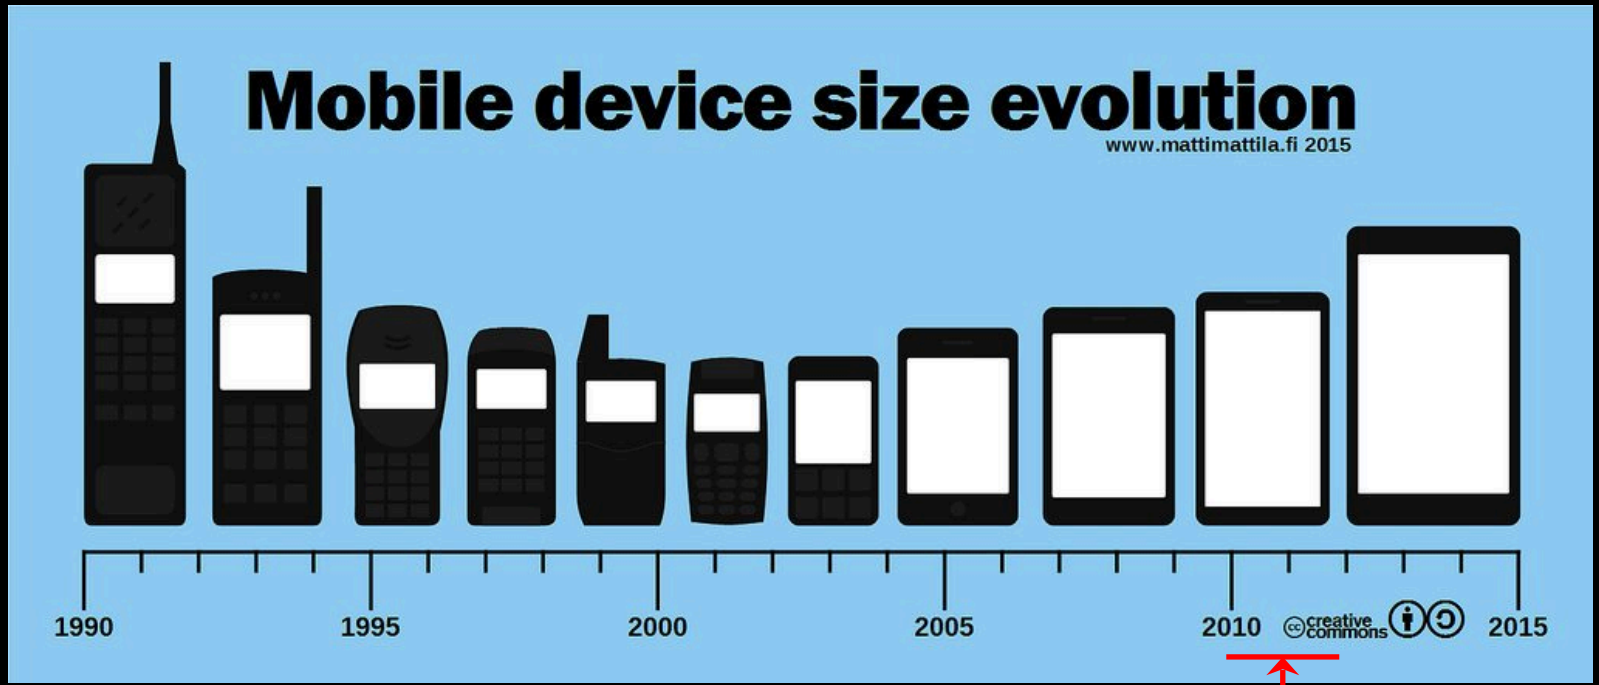
Drastic change in the business paradigm

- Overabundant content: worthless!
- Super-abundant content: search is crucial!
- How to profit from search? Advertising!
- How to profit from social networks? Advertising!
- Online commerce: Recommendation!



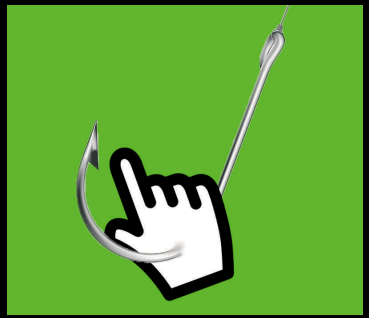
The key technology of all this: AI!

Technological revolutions



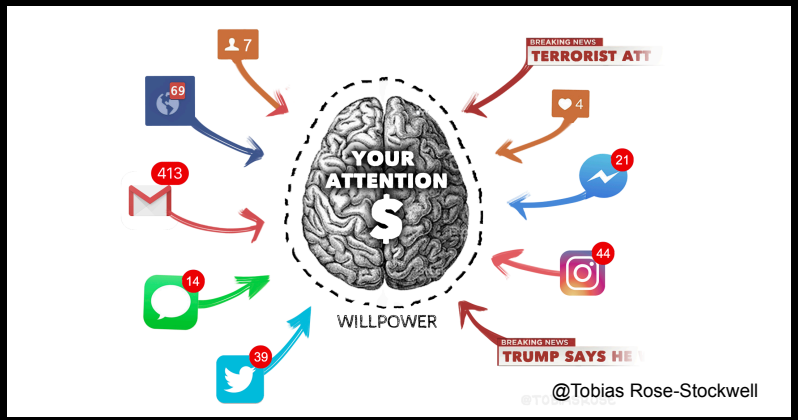
Mobile solutions in social media

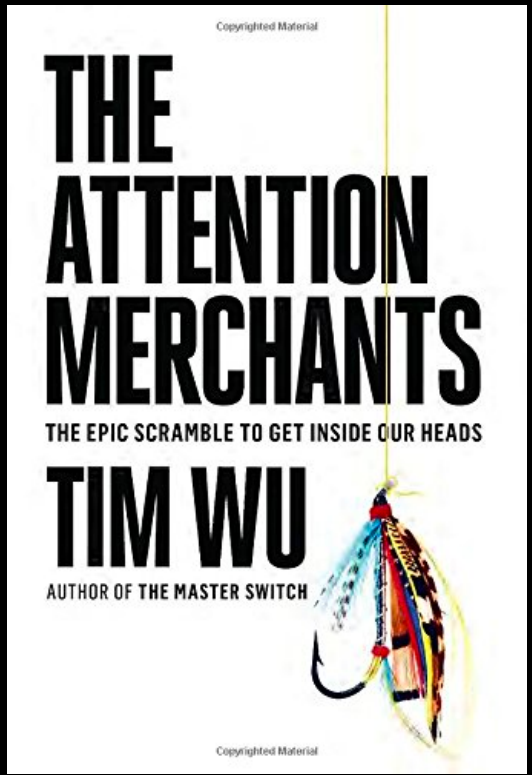


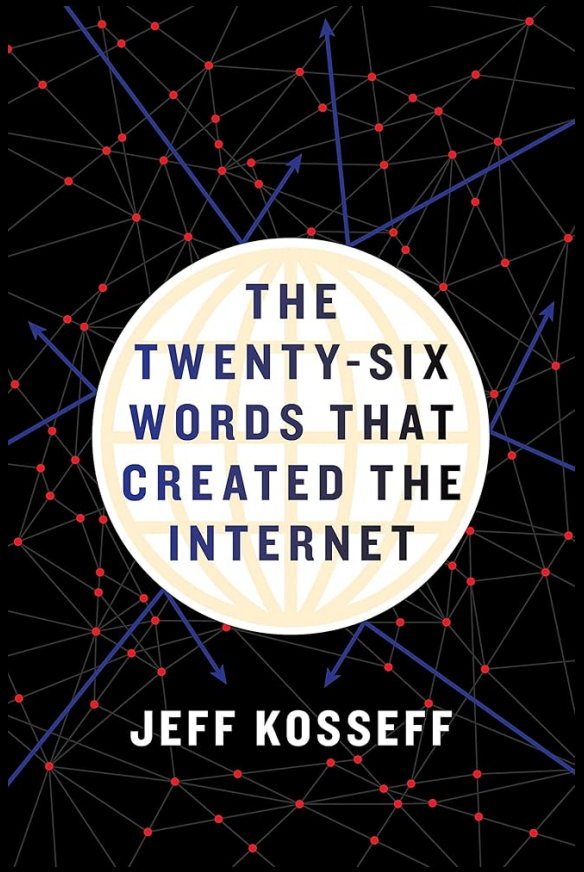


Capture and sell attention, preferences, information about:

tastes, habits, friends, places, activities, preferences (political, sexual, religious, ...)

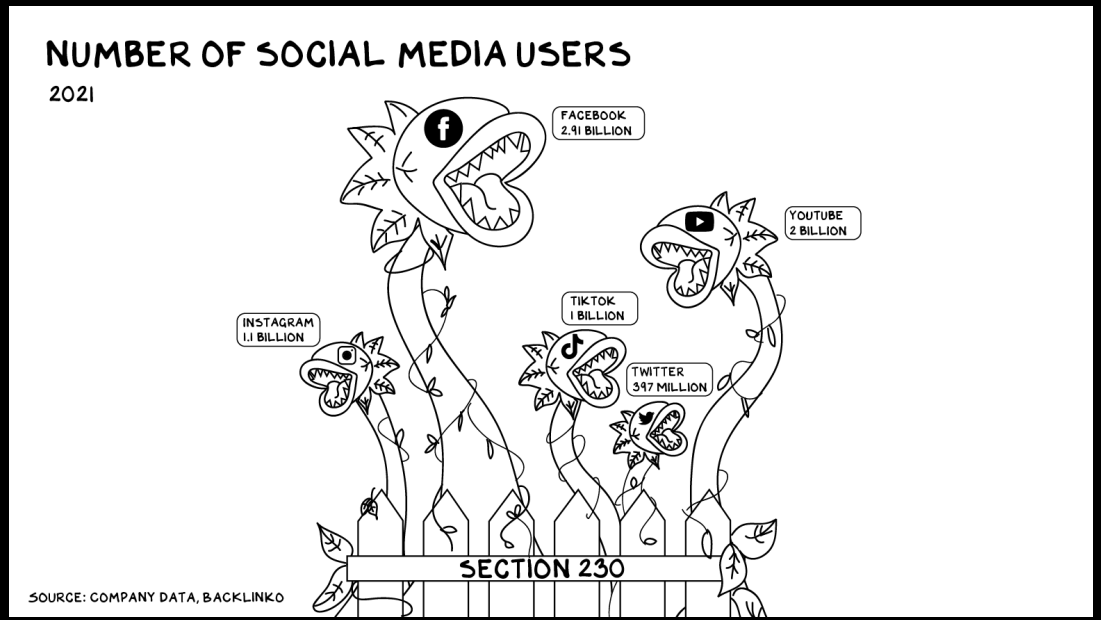






> Section 230, a Provision of the Communication Decency Act, states that, "No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider" (47 U.S.C. § 230). This means is that social media companies are not on the hook for the content published by its users.

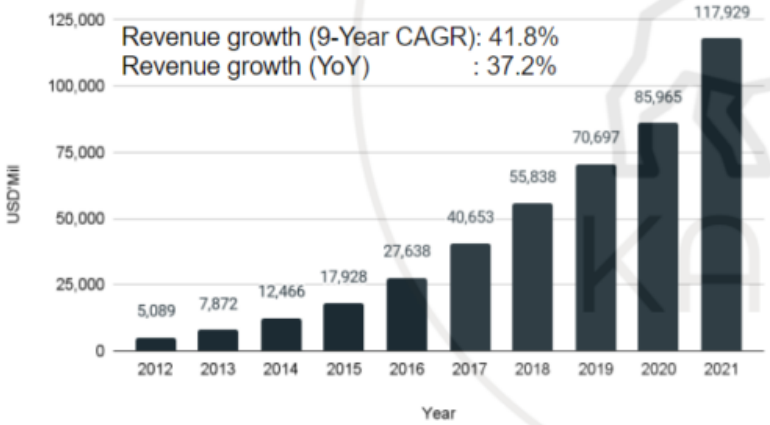
1996



Meta

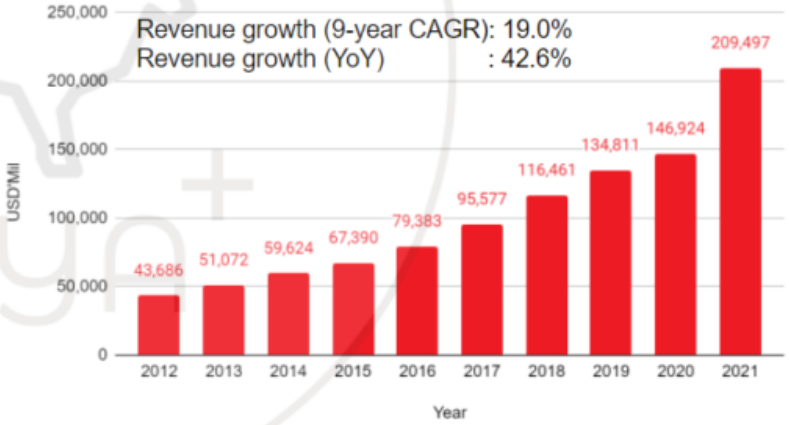
Alphabet

Meta Platforms Inc, Historical advertising revenue



Note: Reality Labs' revenue contribution for FY 2021 is just 1.92% of total revenue, negligible versus advertising revenue for Meta Platforms Inc.

Alphabet Inc. Historical advertising revenue



Note: Google Cloud & Google other revenue removed. Google advertising consists of Google Search, YouTube ads, Google Network Members' properties

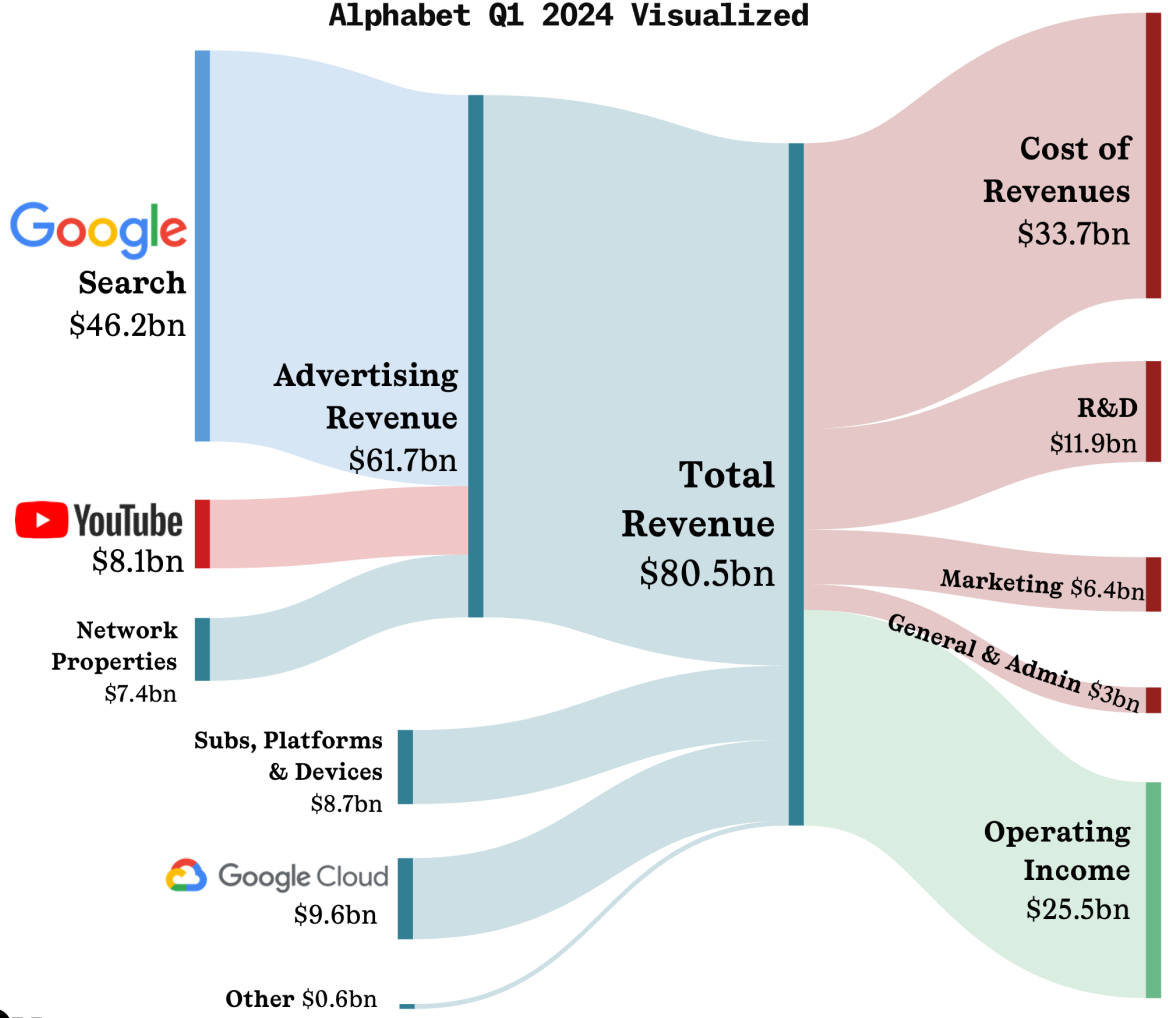
255 B



STILL SEARCHING

Ads Are Still The Profit Engine For Alphabet

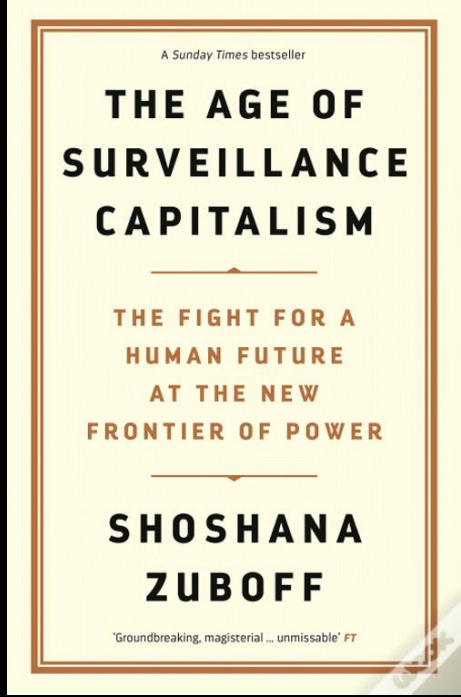
Alphabet Q1 2024 Visualized



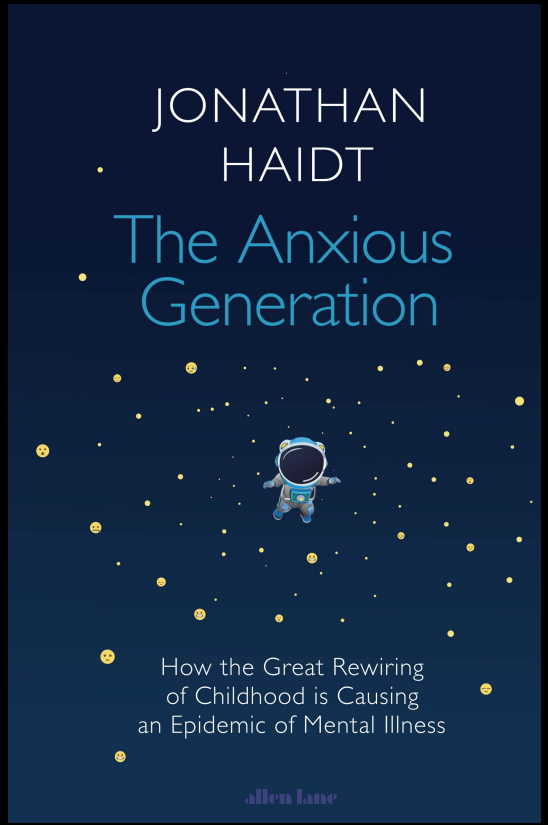
CHARTR

Source: Alphabet

- Business model behind Facebook, Google, Amazon, Twitter, …
- Massive gathering and AI-based analysis of user information.
- Product: behavioral “futures”.
- “If you’re using the service for free, you’re the product.”
- Implications: privacy, addiction, behavior manipulation, threat to democracy, …



- Solutions: regulation, awareness, education, literacy, …



1-The "great rewiring" of childhood: a shift from a **play-based childhood** to a **phone-based childhood**
2-Adults made two opposite mistakes at once: **overprotected** and **underprotected** children

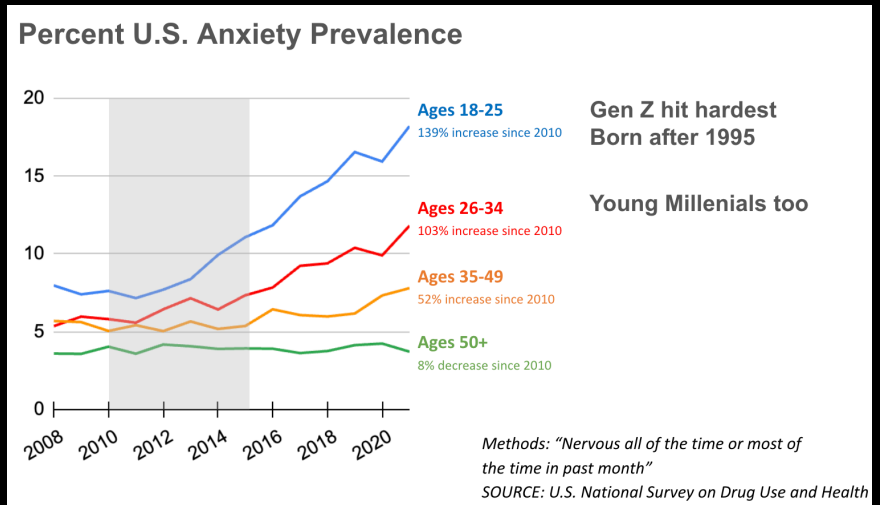
Social Media Is Harming Young People
at a Scale Large Enough to Cause Changes at the Population Level

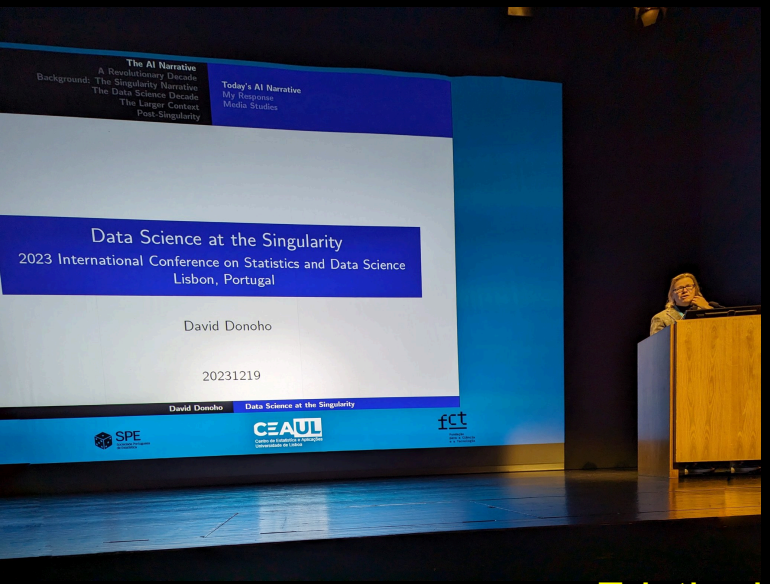
Essay to Be Published in *World Happiness Report 2026*

Jonathan Haidt and Zachary Rausch
New York University Stern School of Business

Corresponding author:
Zachary Rausch (rausch.zachary@nyu.edu)

January 15, 2026





arXiv > stat > arXiv:2310.00865

Statistics > Other Statistics

[Submitted on 2 Oct 2023]

Data Science at the Singularity

David Donoho

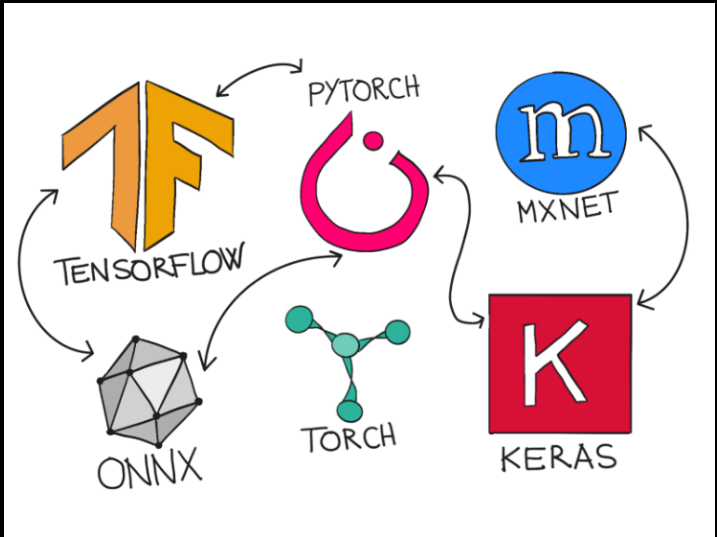
Frictionless reproducibility

[**FR-1: Data**] datafication of everything, with a culture of research data sharing. One can now find datasets publicly available online on a bewildering variety of topics, from chest x-rays to cosmic microwave background measurements to uber routes to geospatial crop identifications.

[**FR-2: Re-execution**] research code sharing including the ability to exactly re-execute the same complete workflow by different researchers.

[**FR-3: Challenges**] adopting challenge problems as a new paradigm powering scientific research. The paradigm includes: a shared public dataset, a prescribed and quantified task performance metric, a set of enrolled competitors seeking to outperform each other on the task, and a public leaderboard. Thousands of such challenges with millions of entries have now taken place, across many fields.

Frictionless reproducibility (Donoho, 2023)



GitHub

Hugging Face

Microsoft

aws

Google Cloud Platform

kaggle

Classical paradigm: supervised learning

Fundamental **bottlenecks**:

- data scarcity
- label scarcity (human expertise is very expensive)

Several approaches:

- data augmentation
- synthetic data
- transfer learning

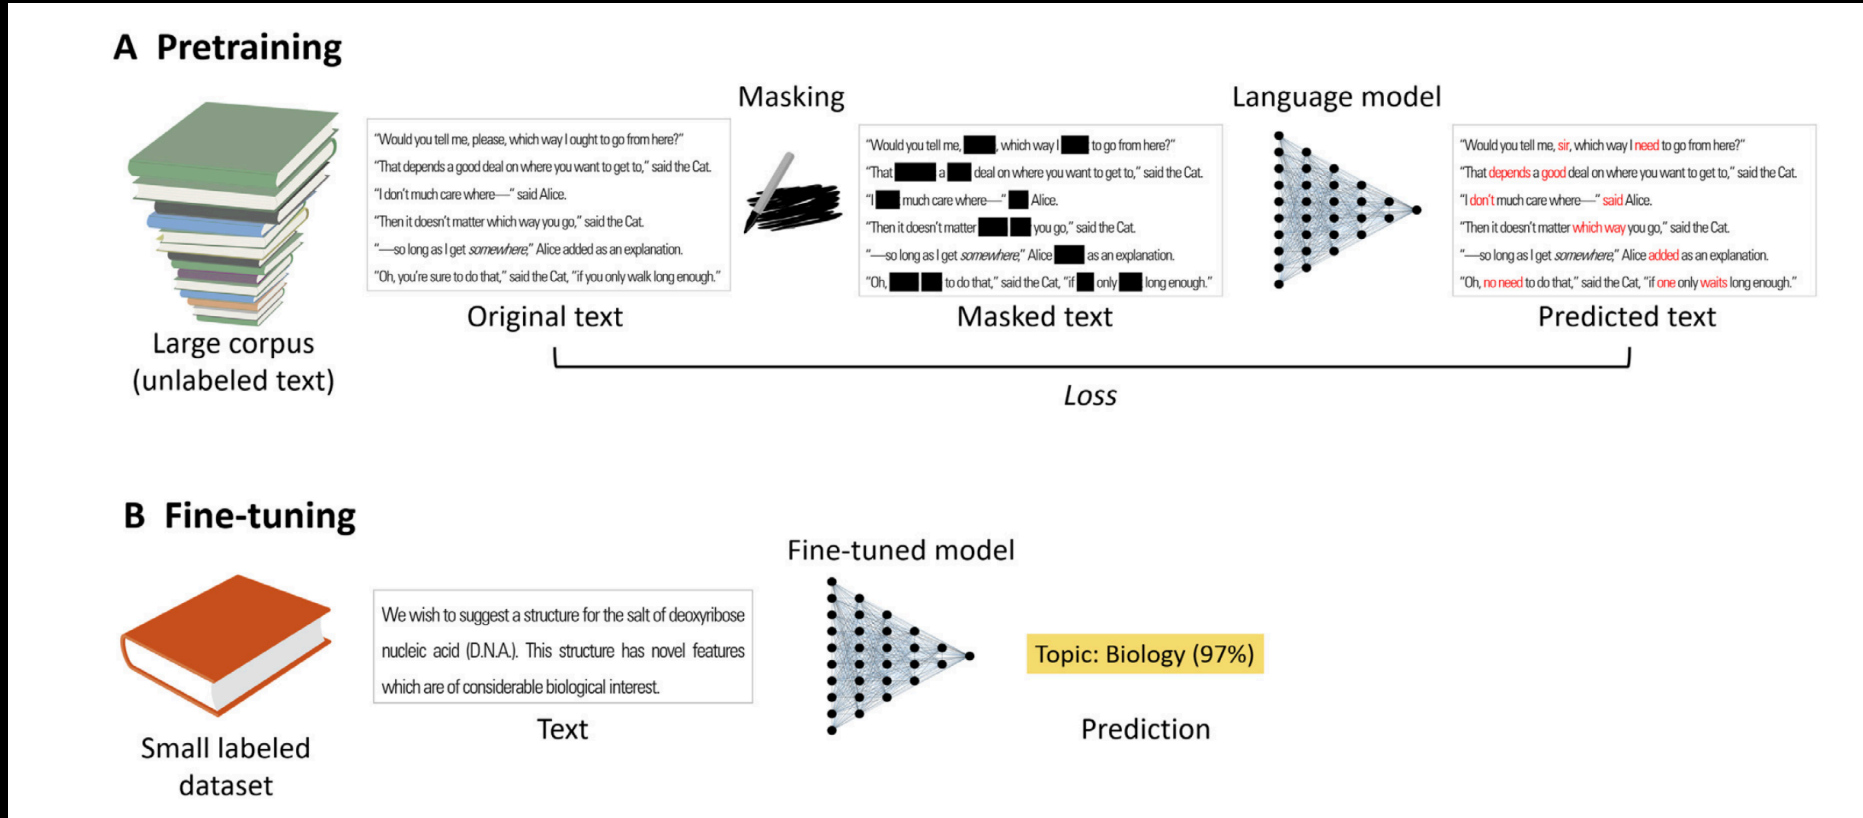
} still supervised learning

- **Paradigm shift: self-supervised learning**

Recent paradigm shift...

Foundation models: pretrain on raw data
(maybe) adapt (fine-tune) to various downstream applications

Self-supervised pretraining: surrogate task, e.g., masked/next word prediction
key advantage: no human intervention!

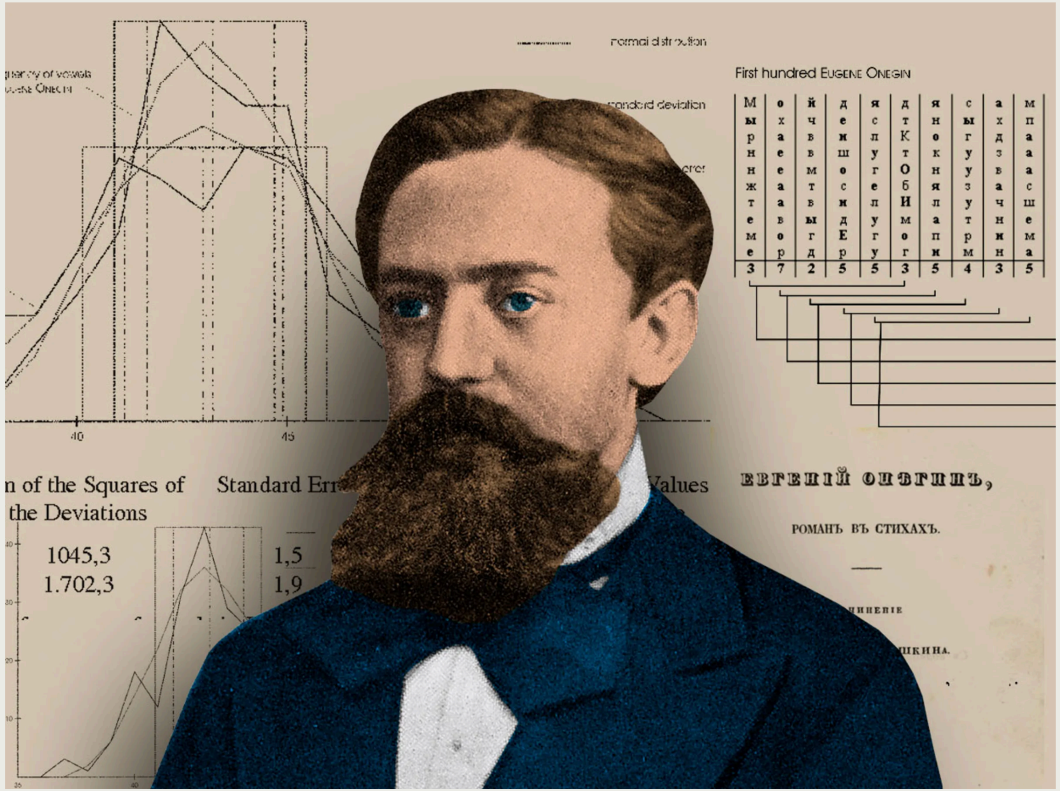


ARTICLE | ARTIFICIAL INTELLIGENCE

Andrey Markov & Claude Shannon Counted Letters to Build the First Language-Generation Models ›

Shannon's said: "OCRO HLI RGWR NMIELWIS"

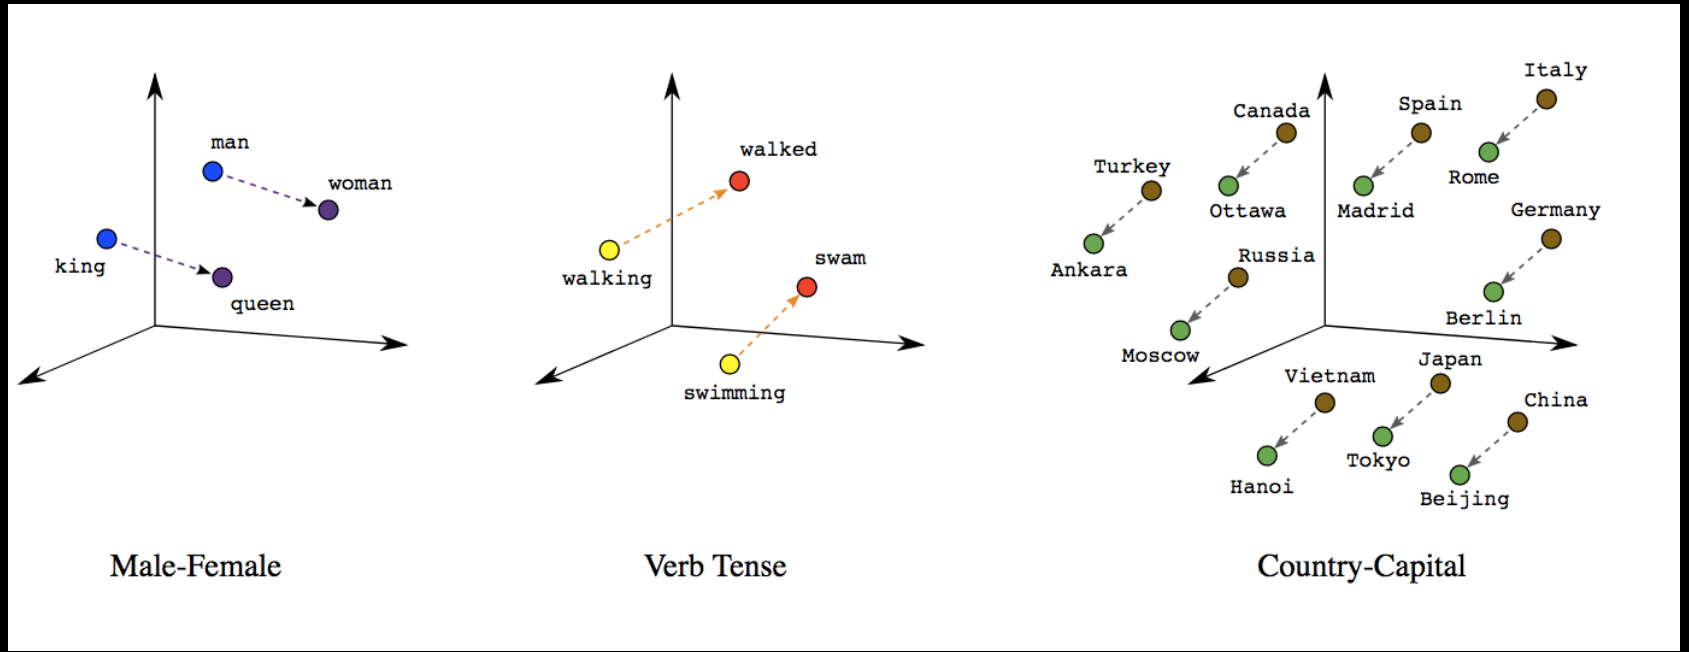
BY OSCAR SCHWARTZ | 11 NOV 2019 | 4 MIN READ |



Russian mathematician Andrey Andreyevich Markov in front of his statistical analyses of Alexander Pushkin's novel, Eugene Onegin.

PHOTO-ILLUSTRATION: GLUEKIT

“Secret sauce” of language models: word embeddings



“Secret sauce” of modern machine learning: embed everything!!!

Attention Is All You Need

Ashish Vaswani* Google Brain avaswani@google.com
Noam Shazeer* Google Brain noam@google.com
Niki Parmar* Google Research nikip@google.com
Jakob Uszkoreit* Google Research usz@google.com
Llion Jones* Google Research llion@google.com
Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu
Łukasz Kaiser* Google Brain lukaszkaiser@google.com
Illia Polosukhin* ‡ illia.polosukhin@gmail.com

Abstract

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.8 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature. We show that the Transformer generalizes well to other tasks by applying it successfully to English constituency parsing both with large and limited training data.

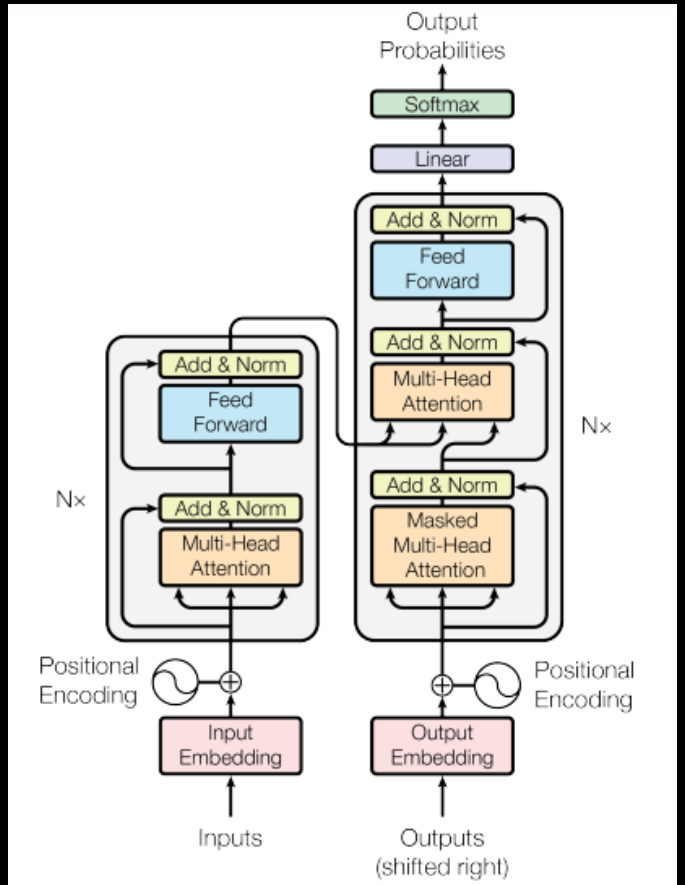
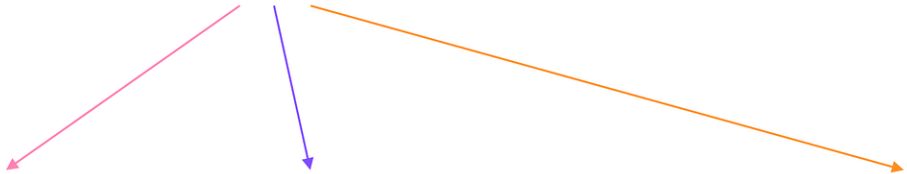


Figure 1: The Transformer - model architecture.

2017

What does **Generative Pre-trained Transformer (GPT)** mean



Generative

Means “**next word prediction.**”

Pre-trained

The LLM is pretrained on massive amounts of text from the internet and other sources.

Transformer

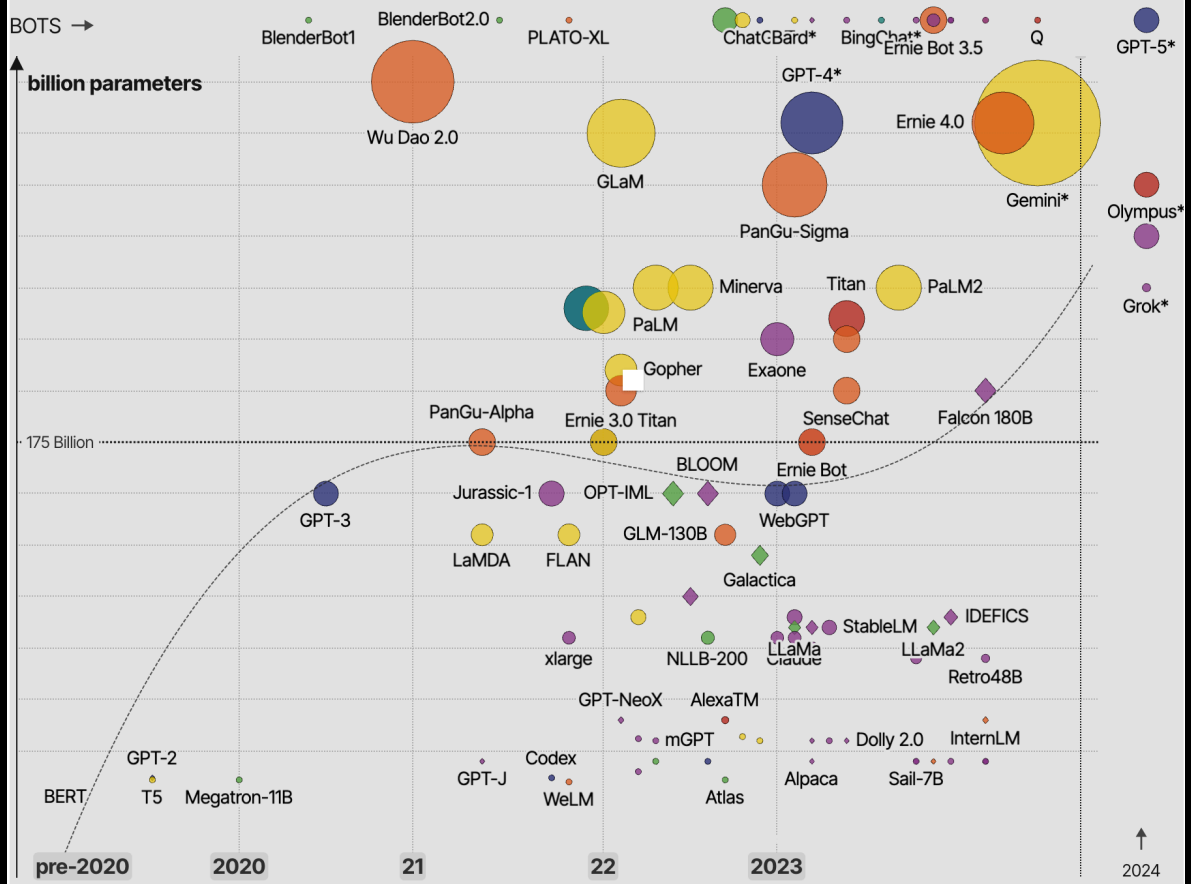
The neural network architecture used (introduced in 2017).

Won’t go into more details here.

The Rise and Rise of A.I.
Large Language Models (LLMs) & their associated bots like ChatGPT

size = no. of parameters open-access

Amazon-owned Chinese Google Meta / Facebook Microsoft OpenAI Other



David McCandless, Tom Evans, Paul Barton
Information is Beautiful // UPDATED 6th Dec 23

source: news reports, LifeArchitect.ai
* = parameters undisclosed // see the data

MADE WITH VIZsweet

A blissful fairytale or a dreadful nightmare?

AI in science and society

The Fairytale

- **Protein structures at scale**
  AlphaFold, Nobel Prize 2024
- **Faster discovery**
  drugs, materials, catalysts
- **Patterns in huge data**
  astronomy, genomics, climate
- **Routine work automated**
  lab, analysis, coding
- **Expertise for everyone**
  tutors and assistants

The Nightmare

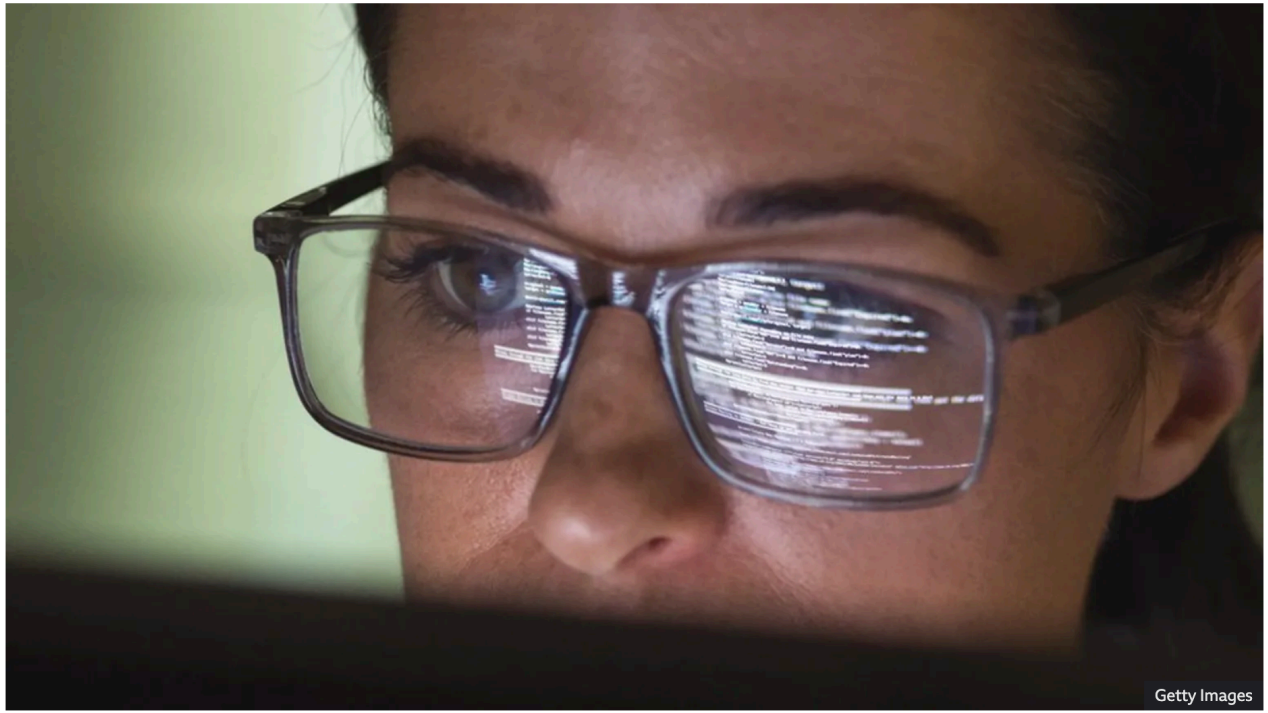
- **Job market disruption**
  and erosion of human skills
- **Loss of control**
  over powerful, opaque systems
- **Unreliable science**
  hallucinations, irreproducibility
- **Misuse**
  deepfakes, surveillance, bio & cyber
- **Concentrated power**
  and a growing energy footprint

AI to hit 40% of jobs and worsen inequality, IMF says

15 January 2024

Share

By **Annabelle Liang,** Business reporter



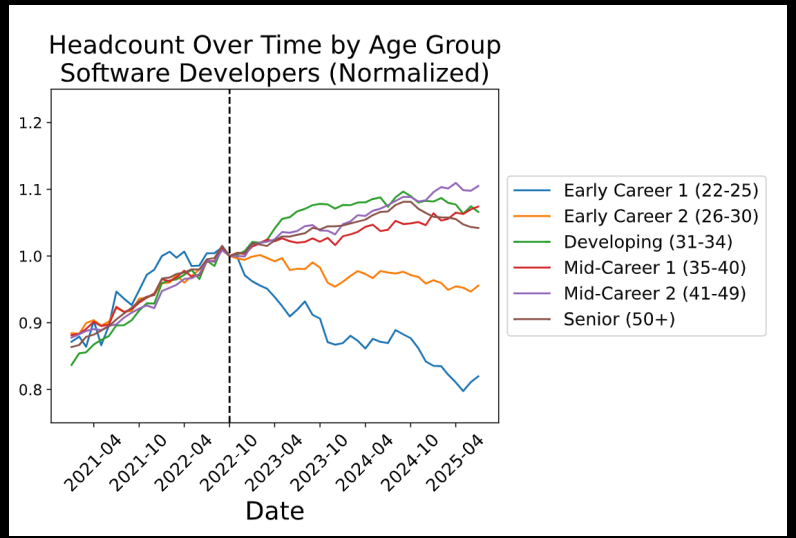
Getty Images

Artificial intelligence is set to affect nearly 40% of all jobs, according to a new analysis by the International Monetary Fund (IMF).

Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence

Erik Brynjolfsson* Bharat Chandar† Ruyu Chen‡§¶

August 26, 2025



THE WALL STREET JOURNAL.

JOANNA STERN

Social-Media Algorithms Rule How We See the World. Good Luck Trying to Stop Them.

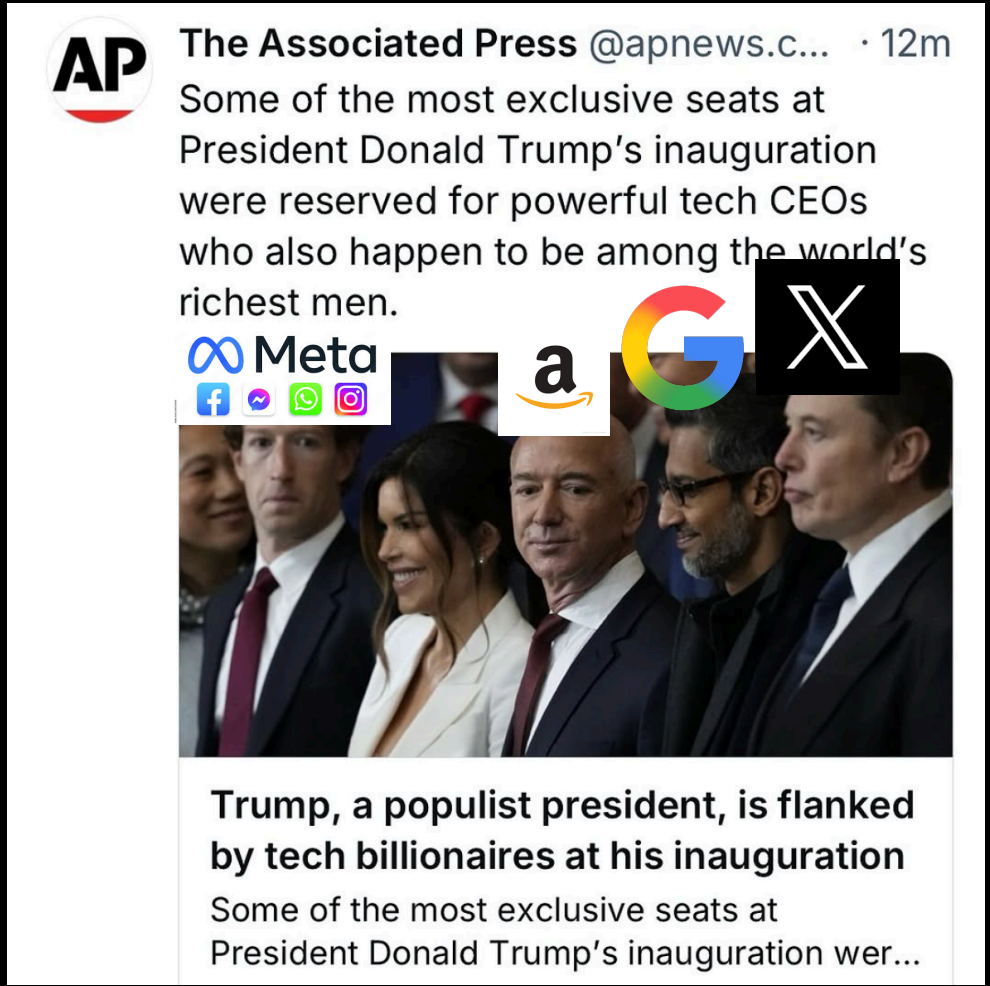
What you see in your feeds isn't up to you. What's at stake is no longer just missing a birthday. It's your sanity—and world peace.



By Joanna Stern Follow

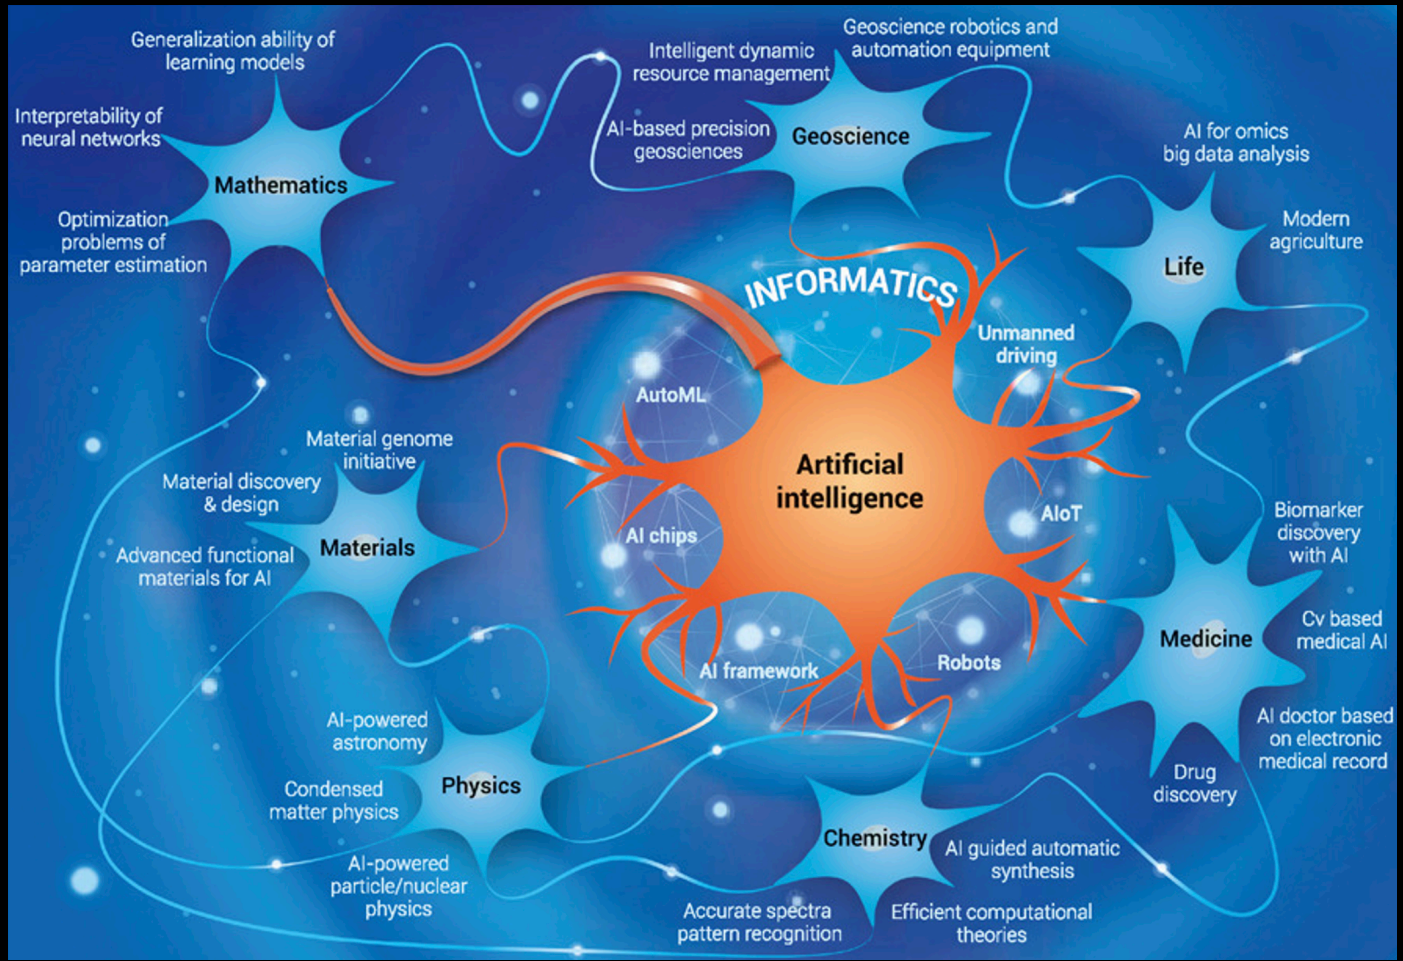
Jan. 17, 2021 7:00 am ET





Is Digital Media a threat to our political systems?

How about scientific impacts?



Forbes

AI

AlphaFold Is The Most Important Achievement In AI—Ever

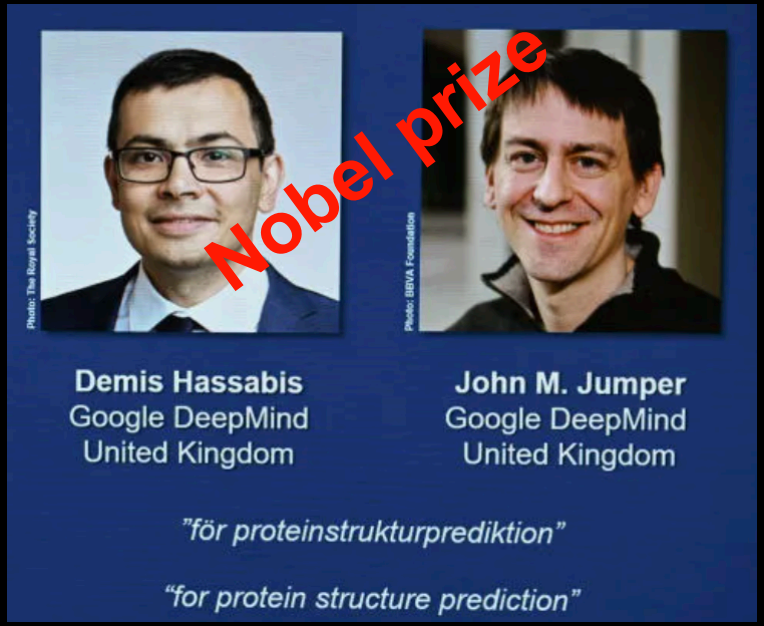
Rob Toews Contributor ⓘ
I write about the big picture of artificial intelligence.

Follow

Oct 3, 2021, 07:34pm EDT



DeepMind's AlphaFold represents the first time a significant scientific problem has been solved by AI.



"This will change medicine. It will change research. It will change bioengineering. It will change everything."

Scientists discover the first new antibiotics in over 60 years using AI

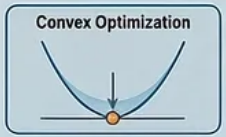


By Oceane Duboust

Published on 31/12/2023 - 16:07 • Updated 16:08

Accelerating Science: How GPT-5 is Changing Research

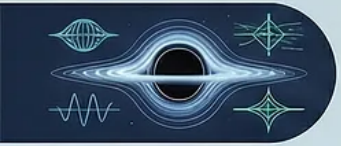
1. Independent Rediscovery of Known Results



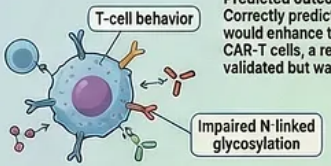
Improved a recent result in Convex Optimization
Improved a key condition from $\eta \in 1/L$ to $\eta \in 1.5/L$, getting halfway to the optimal bound of 1.75/L found by human experts.



Predicted outcomes of unpublished experiments
Correctly predicted that a specific treatment would enhance the cancer-billing ability of CAR-T cells, a result that had been internally validated but was unpublished.



Discovered hidden symmetries of Black Holes
Correctly re-deriving the complex SL(2,R) symmetries of a wave equation on a Kerr background after a "warm-up" on a simpler version.



Uncovered the mechanism behind an immune response
Proposed a non-obvious mechanism (impaired N-linked glycosylation) to explain T-cell behavior, a hypothesis that had eluded lab experts.

2. Deep Literature Search & Conceptual Bridging



Seived 10 "open" Erdős problems by finding forgotten papers
Located docades-old, hard-to-find solutions, navigating cague titles and conceptually related papers with different terminology.



Example: Finding a needle in a 700-citation haystack
Located the correct reference inside a 100-page paper that had over 700 citations for Brdős Problem #339.



Connected new math to existing computer science concepts
identified that a new mathematical result on "α-rolio covers" was conceptually equivalent to "ε-approximate Pareto sets" from computer science.

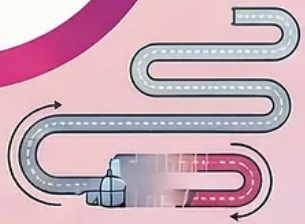


3. Working in Tandem with AI as a Research Partner

Compressed a 6-month calculation into 40 minutes
GPT-5 Pro solved the remaining half of a problem on gravitational rediation from cosmic strings using a more elegant method.



Functioning as a high-speed mathematical collaborator
Mathematician Timothy Gowers ases GPT-5 to quickly test filess, solve subproblems, and find flaws in proof attempts.



Compressed a 6-month calculation into 40 minutes. GPT-5 Pro solved the remaining half of a problem on grovitational radiation from cosmic strings using a more elegant method.



"It made me a one-person army of experts".
Bolit a complex thennonuclear burn propagation model from idea to cimulation to theoretical validation in just 6 hours.



An effective workflow: Human guidance and AI execution
Experts found success by providing high-level sceffelding complex problems with simpler warm-ups, and critically verifying the AI's output.

4. Generating Novel Scientific Results

Solved a genuinely open Erdée problem
Provided the nay creative step that, combined with human insights, iod to the first complete solution for Erdős Problem A648 on squarefree numbers.





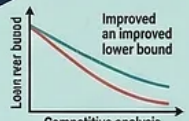


$$29Y - 42P - 1445 \leq K$$

Proved a new open conjecture in graph theory
Developed a new, elegant proof fer a conjoclored loarplelity an anograph counts in treex, a problem felt open in a 2016 paper.



Solved an open problem from the Conference on Learning Theory (COLT).
Resolveo a 2013 open problem by devising a method to disority a lmition parameter in a dscarvic network model based only on a single snapshot.



Improved lower bounds for an online algorithms problem
Created a new obvious counterexample to refute a natural algorithm and provided a key technical idea that helped raise the competitive ratio lower bound.

NotebookLM

Bubeck et al, 2025

IEEE Spectrum

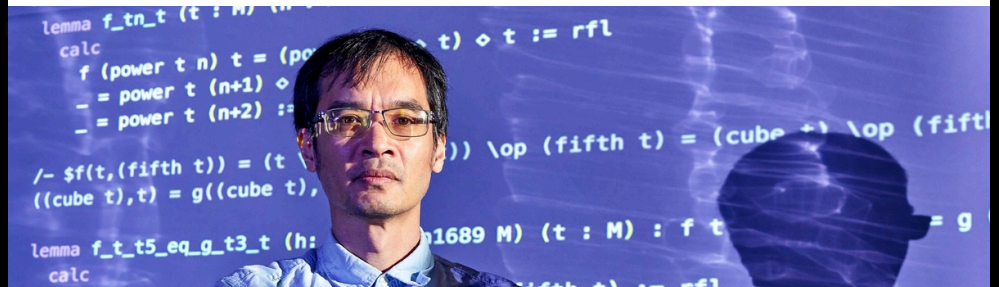
FEATURE AI

What It Means to Be a Mathematician When AI Does the Math

Researchers debate motivation, purpose, and the field's future

BY BENJAMIN SKUSE

25 JUN 2026



Gros Merci!
Many Thanks!
Grande Obrigado!